Lenovo Flex System SI4091 10Gb System Interconnect Module

Application Guide

for Lenovo Enterprise Network Operating System 8.4

Lenovo™

Note: Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices and User Guide* documents on the Lenovo *Documentation* CD and the *Warranty Information* document that comes with the product.

Fourth Edition (October 2017)

Contents

Preface

The *Lenovo Flex System Fabric SI4091 Application Guide for Lenovo Enterprise Network Operating System 8.4* describes how to configure and use the Enterprise NOS 8.4 software on the Lenovo Flex System SI4091 10Gb System Interconnect Module (referred to as SI4091 throughout this document).

For documentation about installing the device physically, see the *Installation Guide* for your SI4091.

Who Should Use This Guide

This guide is intended for network installers and system administrators engaged in configuring and maintaining a network. The administrator should be familiar with Ethernet concepts, IP addressing and SNMP configuration parameters.

What You'll Find in This Guide

This guide will help you plan, implement, and administer Enterprise NOS software. Where possible, each section provides feature overviews, usage examples, and configuration instructions. The following material is included:

Part 1: Getting Started

- Chapter 1, "Introduction," describes basic concepts and workflow for SI4091 usage.
- Chapter 2, "Administrative Access," describes how to access the SI4091 in order to configure the device and view its information. This chapter discusses a variety of manual administration interfaces, including local management via the Console port, and remote administration via Telnet or SNMP.
- Chapter 3, "Initial Setup," describes how to use the built-in Setup utility to perform first-time configuration.
- Chapter 4, "Updating the System Software," describes how to update the software that controls SI4091 operation.

Part 2: Securing the SI4091

- Chapter 5, "Administrative Security," describes methods for changing the default system passwords, using Secure Shell and Secure Copy for administration connections, configuring end-user access control, and placing the SI4091 in protected mode.
- Chapter 6, "Authentication & Authorization Protocols," describes different methods of secure administration for remote administrators. This includes Remote Authentication Dial-in User Service (RADIUS), as well as TACACS+ and LDAP.
- Chapter 7, "Access Control Lists," describes how to use filters to permit or deny specific types of traffic based on a variety of source, destination, and packet attributes.

Part 3: Basic Features

- Chapter 8, "External Uplink Interfaces," describes how to configure loop-free domains using external uplink interfaces.
- Chapter 9, "VLANs," describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs. This chapter also describes Protocol-based VLANs, and Private VLANs.
- Chapter 10, "Ports and Link Aggregation," describes how to group multiple physical ports together to aggregate their bandwidth between large-scale devices.
- Chapter 11, "Quality of Service," discusses Quality of Service (QoS) features, including IP filtering using Access Control Lists (ACLs), Differentiated Services, and IEEE 802.1p priority values.

Part 4: IP Routing

- Chapter 12, "Basic IP Addresses," describes how to configure the BOOTP and DHCP Relay.
- Chapter 13, "Internet Protocol Version 6," describes how to configure the SI4091 for IPv6 host management.
- Chapter 14, "Internet Group Management Protocol," describes how to implement IGMP Snooping to conserve bandwidth in a multicast environment.
- Chapter 15, "Service Location Protocol," describes the Service Location Protocol (SLP) that allows the SI4091 to provide dynamic directory services.

Part 5: Extended Features

- Chapter 16, "CEE and FCoE," discusses using various Converged Enhanced Ethernet (CEE) features such as Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and FIP Snooping for solutions such as Fibre Channel over Ethernet (FCoE).

Part 6: High Availability Fundamentals

- Chapter 17, "Layer 2 Failover," describes how the SI4091 supports high-availability network topologies using Layer 2 Failover with Automatic or Manual Monitoring.
- Chapter 18, "Hot Links," described basic connection redundancy using Hot Links.

Part 7: Network Management

- Chapter 19, "Link Layer Discovery Protocol," describes how Link Layer Discovery Protocol helps neighboring network devices learn about each others' ports and capabilities.
- Chapter 20, "Simple Network Management Protocol," describes how to configure the SI4091 for management through an SNMP client.

Part 8: Appendices

- Appendix A, "Getting Help and Technical Assistance," describes how to get help.
- Appendix B, "Notices," provides trademark and other compliance information.

Additional References

Additional information about installing and configuring the SI4091 is available in the following guides:

- *SI4091 System Interconnect Module User's Guide* (Installation)
- *SI4091 System Interconnect Module Command Reference for Enterprise NOS 8.4*

Typographic Conventions

The following table describes the typographic styles used in this document.

Table 1. *Typographic Conventions*

Typeface or Symbol	Meaning	Example
`ABC123`	This type is used for names of commands, files, and directories used within the text.	View the `readme.txt` file.
	It also depicts on-screen computer output and prompts.	`host#`
`ABC123`	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	`host# sys`
<ABC123>	This italicized type appears in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.	To establish a Telnet session, enter: `host# telnet` *<IP address>*
	This also shows book titles, special terms, or words to be emphasized.	Read your *User's Guide* thoroughly.
`[ ]`	Command items shown inside squared brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	`host# ls [-a]`
`{ | }`	The curled braces and vertical bar are used in command examples to separate choices where multiple options exist. Select only one of the listed options. Do not type the braces or vertical bar.	`host# set {left|right}`
AaBbCc123	This block type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the **Save** button.

Part 1: Getting Started

Chapter 1. Introduction

The Lenovo Flex System SI4091 10Gb System Interconnect Module (referred to as SI4091 throughout this document) provides simplified interconnect options for the Lenovo Flex System and Lenovo PureFlex System environments. The SI4091 facilitates low-latency, media-speed server-to-server traffic within the chassis, based on established Ethernet Layer 2 bridging protocols.

Feature Summary

The SI4091 has the following features:

- Easy Connectivity

  The SI4091 provides a low-touch, easy-to-use connection between server elements within the chassis and the upstream infrastructure. This interface eliminates the complexity normally associated with embedded Layer 2 and Layer 3 switches.

  After initial start-up, the switch uplinks are pre-configured by default to form a single, loop-free IEEE 802.3ad Link Aggregation Group (LAG) that allows for plug-and-play operation. In this state, you can attach as many of the available uplink ports as required to meet application bandwidth requirements to a single upstream device configured with appropriate LACP properties (see "Link Aggregation Control Protocol" on page 151).

- Easy Management

  In addition to the built-in, industry-standard command-line interface, the Lenovo Switch Center tool can be used for centralized management for multiple SI4091 devices:

  - Examining and modifying the configuration of the SI4091.
  - Archiving configurations for faster deployment to new chassis or to replacement SI4091 devices.
  - Continuous monitoring of SI4091 operation, with alerts of events that may impact connectivity.
  - Retrieving traffic statistics and other operational information.

  For management details, see "Administrative Access" on page 25, as well as the *SI4091 System Interconnect Module Command Reference for ENOS 8.4.*

- Loop-Free Operation Without STP

  Unlike traditional switches, the SI4091 prevents data loops among connected links without relying on slow, complicated IEEE 802.1d Spanning-Tree Protocols.

- Server Failover

  To assist in providing seamless failover in the event of connectivity disruptions outside the chassis, Layer 2 Failover is enabled by default. If the number of operational links between the SI4091 and the upstream device falls below a configurable threshold, the SI4091 will close all internal server ports associated with the configured failover trigger, signalling the NICs on the affected servers to initiate failover to an alternate SI4091.

  Using the default configuration, at least one uplink connection must be operational for the internal server ports to remain operational.

  For more information, see Chapter 17, "Layer 2 Failover."

- Converged Data and Storage Applications

  The SI4091 provides convergence for traditional Ethernet traffic and Fibre Channel over Ethernet (FCoE) storage session traffic on all internal and external links, thus reducing the need for separate data and Storage Area Network (SAN) infrastructures.

  For more information, see "CEE and FCoE" on page 197.

Recommended System Deployments

The following are recommended SI4091 deployments.

Basic Topology

One recommended SI4091 deployment is shown in Figure 1:

Figure 1. Basic SI4091 Topology

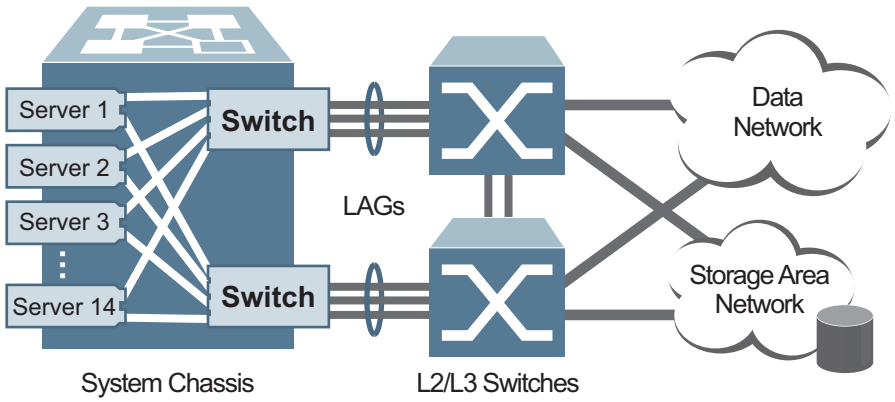

In this deployment, a pair of SI4091 devices is used in each Lenovo Flex System chassis, installed in I/O Bays 1 and 2, with each assigned a unique IP address and default gateway via the Flex Chassis Management Module (CMM).

This provides redundant data paths through the chassis while also providing additional internal and external bandwidth during normal operation. Additional SI4091 devices can be installed in I/O Bays 3 and 4 when additional I/O bandwidth or server NICs are needed. In addition, each SI4091 device is attached to separate external switches to achieve the desired redundancy, protecting against a single point of failure.

Each server has a dual-port 10 Gbps Ethernet interface, with each interface independently connected to one of the two SI4091 devices. For each server, optional 4-port and 8-port mezzanine cards are available for applications that require greater than 20 Gbps of bandwidth per server, or additional redundancy.

It is recommended that NIC teaming (or bonding) be configured on each server. NIC teaming allows redundant NICs on each server to create an active path and a backup path to the pair of SI4091 devices. Loss of one SI4091 device (such as when it is removed from the chassis), or removal or loss of all uplinks due to failure of an upstream device, will force the corresponding internal server links to close, thus signalling the NIC teaming capability in the server to use the backup path.

The multiple external ports on the SI4091 are grouped together to form a single, loop-free channel (also known as a Link Aggregation Group) to the upstream switch. Two uplinks per SI4091 are the minimum recommended for redundancy.

During SI4091 operation, at least one active uplink port is required for each uplink. Otherwise, the internal ports for the LAG will remain inactive. This avoids situations where an internal server path is active but has no corresponding path to the external domains.

If desired, converged network adapters on each server allow Fibre Channel storage traffic to share the internal and external links with the normal Ethernet traffic. The corresponding FCoE support must be configured in the external upstream device, such as the Lenovo RackSwitch G8264CS, to redirect the Fibre Channel traffic streams to the corresponding Storage Area Network (SAN).

As with NIC teaming, FCoE multi-path protocols on the server will redirect FCoE session traffic to an alternate path when the primary path is disrupted.

VLAG Topology Modification

An alternate topology is shown in Figure 2:

Figure 2. SI4091 VLAG Topology

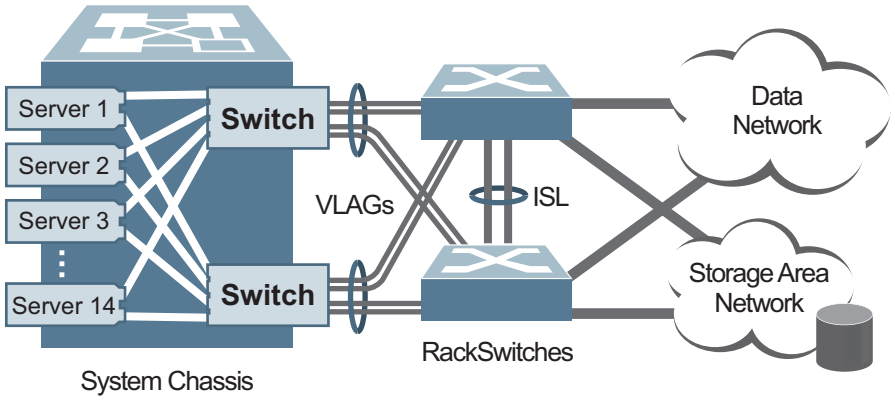

This deployment takes advantage of the Virtual Link Aggregation Group (VLAG) capability available when the SI4091 is connected to the VLAG-capable switches such as the Lenovo RackSwitch G8264 or RackSwitch G8264CS. This allows the SI4091 uplink channel to be physically distributed across a pair of upstream switches. The RackSwitch devices are peered using an inter-switch link (ISL) so that they act in concert, appearing as if a single switch to each SI4091 device.

As with the basic topology, the VLAG topology utilizes NIC teaming and FCoE multi-path redundancy options configured at the server. These options provide seamless failover in the event of connectivity disruptions outside the chassis.

SAN domains, when present, are managed and controlled via the external converged network and storage devices. For example, using a Lenovo RackSwitch G8264CS as the uplink switch can provide converged Ethernet and 8-Gbps Fibre Channel storage access, with the capability to divert the storage session traffic from each chassis to the targeted storage elements.

The SI4091 can also be used to interface to Cisco Nexus 50XX converged network switches.

Recommended Workflow

The SI4091 is ready to provide limited function right out of the box. However, to utilize the device most effectively, some degree of custom configuration is recommended. The most common steps for SI4091 deployment are as follows:

- Access the SI4091 command line (see Chapter 2, “).
- Use the Setup wizard to initialize the system (see Chapter 3, “).
- Change default passwords (see Chapter 5, “).
- Update firmware to the most current version (see Chapter 4, “).
- Configure uplink interfaces (see Chapter 8, “External Uplink Interfaces) and VLANs (see Chapter 9, “).
- Configure optional FCoE if desired (see Chapter 16, “).
- Configure other optional features.
- Back-up your configuration.
- Monitor and adjust operations as necessary.

Note: The steps listed are for a generic deployment. Depending on your specific requirements, your actual workflow may differ.

Chapter 2. Administrative Access

The SI4091 is ready to perform basic functions right out of the box. Some of the more advanced features, however, require some administrative set-up before they can be used effectively.

The SI4091 provides a variety of options for accessing the device to perform configuration, and to view operational information and statistics.

This chapter discusses the various methods that can be used to administer the SI4091.

Administration Interfaces

The SI4091 software provides a variety of user-interfaces for administration:

- The Flex System Chassis Management Module (CMM). The Flex System chassis includes a CMM as the central element for overall chassis management and control. Using the tools available through the CMM, the administrator can configure many of the SI4091 features and can also access other SI4091 administration interfaces.
- A built-in, text-based command-line interface is available for access via the Console port or an optional Telnet or SSH session.
- SNMP support for access through network management software.

The specific interface chosen for an administrative session depends on user preferences, as well as the SI4091 configuration and the available client tools.

In all cases, administration requires that the SI4091 hardware is properly installed and turned on (see the *SI4091 System Interconnect Module User's Guide*).

Establishing a Connection

The factory default settings permit initial switch administration through the built-in serial port, as well as default IP addresses on VLAN 1 and the out-of-band management port.

To facilitate switch access, the in-band and out-of-band management interfaces are configured with factory default IP addresses, as follows:

- VLAN 1/Interface 1: `192.168.49.50/24`
- Out-of-band Management Port 1: `192.168.50.50/24`

Remote access using the network requires the accessing terminal to have a valid, routable connection to the SI4091 interface. This requires that the SI4091 be configured with a client IP address. The address may be configured manually, or an IPv4 address can be provided automatically by the SI4091 using a service such as DHCP or BOOTP client (see "BOOTP/DHCP Client IP Address Services" on page 31), or an IPv6 address can be obtained using IPv6 stateless address configuration.

Note: Throughout this manual, *IP address* is used in places where either an IPv4 or IPv6 address is allowed. IPv4 addresses are entered in dotted-decimal notation (for example, 10.10.10.1), while IPv6 addresses are entered in hexadecimal notation (for example, 2001:db8:85a3::8a2e:370:7334). In places where only one type of address is allowed, *IPv4 address* or *IPv6 address* is specified.

Using the Chassis Management Module

The SI4091 is an integral subsystem within the overall Lenovo Flex System. The Flex System chassis includes a chassis management module (CMM) as the central element for overall chassis management and control.

The SI4091 uses port 26 (MGT1) to communicate with the chassis management module(s). Even when the SI4091 is in a factory default configuration, you can use the 1Gb Ethernet port on each CMM to configure and manage the SI4091.

For more information about using the chassis management module, see the SI4091 installation guide.

Factory-Default vs. CMM-Assigned IP Addresses

Each SI4091 must be assigned its own Internet Protocol version 4 (IPv4) address, which is used for communication with an SNMP network manager or other transmission control protocol/Internet Protocol (TCP/IP) applications (for example, BOOTP or TFTP). The factory-default IPv4 address is 10.90.90.9*x*, where *x* is based on the number of the bay into which the SI4091 is installed. For additional information, see the *Installation Guide*. The chassis management module assigns an IPv4 address of 192.168.70.12*x*, where *x* is also based on the number of the bay into which each SI4091 is installed, as shown in the following table:

Table 2. *SI4091 IPv4 addresses, by chassis module bay numbers*

Bay Number	Factory-Default IPv4 Address	IPv4 Address Assigned by CMM
Bay 1	10.90.90.91	192.168.70.120
Bay 2	10.90.90.92	192.168.70.121
Bay 3	10.90.90.93	192.168.70.122
Bay 4	10.90.90.94	192.168.70.123

Note: SI4091s installed in Bay 1 and Bay 2 connect to server NICs 1 and 2, respectively.

Using Telnet

A Telnet connection offers the convenience of accessing the SI4091 from a workstation connected to the network. Telnet access provides the same options for user and administrator access as those available through the Console port.

By default, Telnet access is disabled. After initial setup of the SI4091 (see Chapter 3, "Initial Setup"), use the following commands (only available when using the Console port) to enable or disable Telnet access:

```
SI 4091> enable
SI 4091# configure terminal
SI 4091(config)# [no] access telnet enable
```

Once the SI4091 is configured with an IP address and gateway, you can use Telnet to access device administration from any workstation connected to the management network.

To establish a Telnet connection with the SI4091, run the Telnet program on client your workstation and issue the following Telnet command:

```
telnet <SI4091 IPv4 or IPv6 address>
```

You will then be prompted to enter a password as explained "System Login Levels" on page 38.

Using Secure Shell

Although a remote administrator can manage the configuration of a SI4091 via Telnet, this method does not provide a secure connection. The Secure Shell (SSH) protocol enables you to securely log in over a network to execute commands remotely. As a secure alternative to using Telnet to manage SI4091 configuration, SSH ensures that the management session is encrypted.

The SI4091 supports only one session of key/cipher generation at a time. Thus, a SSH/SCP client will not be able to login if the SI4091 is performing key generation at that time. Similarly, the system will fail key generation if an SSH/SCP client is logging in at that time.

The SI4091 supports the following encryption and authentication methods for SSH and SCP:

- Server Host Authentication: 1024-bit RSA host key
- Key Exchange: diffie-hellman-group1-sha1, diffie-hellman-group14-sha1
- Encryption: 3des-cbc, aes128-cbc, aes128-ctr, arcfour, arcfour128, arcfour256, blowfish-cbc, rijndael128-cbc
- MAC: hmac-md5, hmac-md5-96, hmac-sha1, hmac-sha1-96
- User Authentication: Local password authentication, LDAP, RADIUS, TACACS+

Using SSH to Access the SI4091

By default, the SSH feature is enabled. For information about enabling and using SSH for SI4091 access, see "Secure Shell and Secure Copy" on page 69.

Once the IP parameters are configured, you can access the command line interface using an SSH connection.

To establish an SSH connection with the SI4091, run the SSH program on your cleint workstation by issuing the SSH command, followed by the SI4091 IPv4 or IPv6 address:

```
ssh <SI4091 IP address>
```

You will then be prompted to enter a password as explained "System Login Levels" on page 38.

Using Simple Network Management Protocol

SI4091 software provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software.

Note: SNMP is disabled by default. However, if community strings are already configured on the switch, any software update will leave SNMP enabled.

To access the SNMP agent on the SI4091, the read and write community strings on the SNMP manager must be configured to match those on the SI4091. The default read community string on the SI4091 is `public` and the default write community string is `private`.

The read and write community strings on the SI4091 can be changed using the following privileged configuration commands:

```
SI 4091(config)# snmp-server read-community <1-32 characters>
  -and-
SI 4091(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager can to reach any of the IP interfaces on the SI4091.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the SI4091, configure the trap host on the SI4091 with the following commands:

```
SI 4091(config)# snmp-server trap source <trap source IP interface>
SI 4091(config)# snmp-server host <IPv4 address> <trap host community string>
```

For more information on SNMP usage and configuration, see "Simple Network Management Protocol" on page 257.

BOOTP/DHCP Client IP Address Services

For remote SI4091 administration, the client terminal device must have a valid IP address on the same network as an interface on the SI4091. The IP address on the client device may be configured manually, or obtained automatically using IPv6 stateless address configuration, or an IPv4 address may obtained automatically via BOOTP or DHCP relay as discussed below.

When the SI4091 receives a BOOTP/DHCP request from a client seeking an IPv4 address, the SI4091 acts as a proxy for the client. The request is forwarded as a UDP Unicast MAC layer message to the BOOTP/DHCP servers configured for the client's VLAN, or to the global BOOTP/DHCP servers if no domain-specific BOOTP/DHCP servers are configured for the client's VLAN. The servers respond to the SI4091 with a Unicast reply that contains the IPv4 default gateway and the IPv4 address for the client. The SI4091 then forwards this reply back to the client.

DHCP is described in RFC 2131. DHCP uses UDP as its transport protocol. The client sends messages to the server on port 67 and the server sends messages to the client on port 68.

BOOTP and DHCP relay are collectively configured using the BOOTP commands on the SI4091.

For more information, see "Basic IP Addresses" on page 171.

Host Name Configuration

The SI4091 supports DHCP host name configuration as described in RFC 2132, option 12. DHCP host name configuration is enabled by default.

Host name can be manually configured using the following command:

```
SI 4091(config)# hostname <name>
```

If the host name is manually configured, the SI4091 does not replace it with the host name received from the DHCP server.

After the host name is configured on the SI4091, if DHCP or DHCP host name configuration is disabled, the SI4091 retains the host name.

To help avoid misconfiguration during management of multiple similar devices, the SI4091 command prompt includes the host name.

Host name configuration can be enabled/disabled using the following command:

```
SI 4091(config)# [no] system dhcp hostname
```

SYSLOG Server

During SI4091 startup, if the system fails to read its configuration file, a message can be recorded in the SYSLOG server.

The SI4091 supports requesting a SYSLOG server IP address from the DHCP server as described in RFC 2132, option 7. DHCP SYSLOG server request option is enabled by default.

If SYSLOG server address is manually configured, it will take priority over a DHCP-assigned SYSLOG server.

Up to two SYSLOG server addresses received from the DHCP server can be used. The SYSLOG server address can be obtained over a management port or a data port.

Use the `show logging` command to view the SYSLOG server address.

DHCP SYSLOG server address option can be enabled/disabled using the following command:

```
SI 4091(config)# [no] system dhcp syslog
```

Easy Connect Wizard

Lenovo EasyConnect (EZC) is a feature designed to simplify switch configuration. A set of predefined configurations can be applied on the switch via ISCLI. By launching the EZC Wizard, you are prompted for a minimal set of input and the tool automatically customizes the switch software.

The EZC Wizard allows you to choose one of the following configuration modes:

- **Basic System** mode supports settings for hostname, static management port IP, netmask, and gateway.
- **Transparent** mode collects server and uplink port settings. vNIC groups are used to define the loop free domains.

  Note: You can either accept the static defaults or enter a different port list for uplink and/or server ports.
- **Redundant** mode refers to VLAG settings.

The EZC configuration will be applied immediately. Any existing configuration will be deleted, the current active or running configuration will not be merged or appended to the EZC configuration.

For any custom settings that are not included in the predefined configuration sets, the user has to do it manually.

Notes:

- EZC is not available in stacking mode.
- To support scripting, the feature also has a single-line format. For more information, please refer to *Lenovo Networking ISCLI Reference Guide*.

Note: To support scripting, the feature also has a single-line format. For more information, please refer to *Lenovo Networking ISCLI Reference Guide*.

Using the Easy Connect Wizard

To launch the EZC Wizard, use the following command:

```
SI 4091# easyconnect
```

The wizard displays the available predefined configuration modes. You are prompted to select one of the following options:

```
SI 4091# easyconnect
Auto configures the switch into a set configuration based on the input
provided.
Current configuration will be overwritten with auto configuration
settings.
The wizard can be canceled anytime by pressing Ctrl+C.
Select which of the following features you want enabled:
#Configure Basic system (yes/no)?
#Configure Transparent mode (yes/no)?
```

Basic System Mode Configuration Example

This example shows the parameters available for configuration in Basic System mode:

```
SI 4091# easyconnect
Configure Basic system (yes/no)? y

Please enter "none" for no hostname.
Enter hostname(Default: None)? host

Please enter "dhcp" for dhcp IP.
Select management IP address (Current: 10.241.13.32)?
Enter management netmask(Current: 255.255.255.128)?
Enter management gateway:(Current: 10.241.13.1)?

Pending switch port configuration:

    Hostname: host
    Management interface:
        IP:      10.241.13.32
        Netmask: 255.255.255.128
        Gateway: 10.241.13.1
Confirm erasing current config to re-configure Easy Connect (yes/no)?
```

```
SI 4091# easyconnect
Configure Basic system (yes/no)? y

Please enter "none" for no hostname.
Enter hostname(Default: None)? Host

Select management port number(Default: 1)?

Please enter "dhcp" for dhcp IP.
Select management IP address (Current: 10.241.13.32)?
Enter management netmask(Current: 255.255.255.128)?
Enter management gateway:(Current: 10.241.13.1)?

Pending switch port configuration:

    Hostname: host
    Management interface:
        Port:    1
        IP:      10.241.13.32
        Netmask: 255.255.255.128
        Gateway: 10.241.13.1
```

Note: You can either accept the default values or enter new parameters.

Transparent Mode Configuration Example

This example shows the parameters available for configuration in Transparent mode:

```
SI 4091# # easyconnect
Configure Transparent mode (yes/no)? y
Select Uplink Ports (Static Defaults:  17-24)?
The following Uplink ports will be enabled:
        Uplink ports(1G/10G):  17-24
Select Server Ports (Static Defaults:  25-64)?
The following Server ports will be enabled:
        Server ports(1G/10G):  25-64
Pending switch configuration:

    Uplink Ports:    17-24
    Server Ports:    25-64
    Disabled Ports:  1,5,9,13
Confirm erasing current config to re-configure Easy Connect (yes/no)?
```

Notes:

- If your selection for a port group contains ports of different mode or speed, the selection is not valid and you are guided to either select other ports or change the speed of the ports.
- If your selection for an uplink port group contains ports of different mode or speed, the selection is not valid and you are guided to select other ports. Server ports can have ports of different mode or speed selected at the same time.
- You can either accept the static defaults or enter a different port list for uplink and/or server ports.

Redundant Mode Configuration Example

This example shows the parameters available for configuration in Redundant mode:

```
SI 4091# #easyconnect
Configure Switch Redundant mode (yes/no)? y

Note: It is recommended to select Basic system configuration in order to
set the management IP address used for vLAG health check.

Configure Basic system (yes/no)? y

Configure this switch as vLAG Primary or Secondary Peer
(primary/secondary)? prim

The following ISL ports will be enabled:

Select vLAG TierID (Default: 101)?

Select management IP address (Current: 192.168.49.50)?

Enter management netmask (Current: 255.255.255.0)?

Select Peer IP address for vLAG healthcheck (Default: 1.1.1.2)?
Warning: vLAG healthcheck Peer IP is not reachable.
Do you want to select another Peer IP (yes/no)? y
Select Peer IP address for vLAG healthcheck (Default: 1.1.1.2)?
Warning: vLAG healthcheck Peer IP is not reachable.
Do you want to select another Peer IP (yes/no)? n

The following Uplink ports will be enabled:

The following Downlink ports will be enabled:
```

```
Please enter "none" for no hostname.
Enter hostname(Default: Primary VLAG)?

Please enter "none" for no gateway.
Enter management gateway:(Default: 0.0.0.0)?

Pending switch configuration:

    vLAG switch type:   Primary
    ISL Ports:
    vLAG TierID:        101
    vLAG Peer IP:       1.1.1.2
    Uplink Ports:
    Downlink Ports:
    Disabled Ports:      empty

    Hostname: Primary VLAG
 Management interface:
        IP:      192.168.49.50
        Netmask: 255.255.255.0
        Gateway: 0.0.0.0

Confirm erasing current config to re-configure Easy Connect (yes/no)?
```

Notes:

- If your selection for a port group contains ports of different speed, the selection is not valid, and you are guided to either select other ports or change the speed of the ports.
- All unused port are configured as shut down in the configuration dump.
- You can either accept the static defaults or enter a different port list for ISL, uplink, and/or downlink ports.

System Login Levels

To enable better system management and user accountability, three levels or *classes* of system access have been implemented on the SI4091. Privileges for each level of access increase as needed to perform various management tasks. Conceptually, access classes are defined as follows:

- User—Interaction with the SI4091 is completely passive: nothing can be changed on the system. Users may display information that has no security or privacy implications, such as device statistics and current operational state information.
- Operator—Operators can only effect temporary changes on the SI4091. These changes will be lost when the system is rebooted/reset. Operators have access to the system management features used for regular operations. Because any changes an operator makes are undone when the device is reset, operators cannot permanently impact operation.
- Administrators—Administrators are the only ones that may make permanent changes to the SI4091 configuration (those that are persistent across a reboot/reset of the device). Administrator access is used for configuring and troubleshooting the SI4091. Because administrators can also make temporary (operator-level) changes, they must be aware of the interactions between temporary and permanent changes.

Access to the SI4091 is controlled through the use of unique account names and passwords. Once you are connected to the SI4091 via Console port, remote Telnet, or SSH, you are prompted to enter a password. The default account names and password for each access level are listed in the Table 3 on page 38.

Note: It is recommended that you change default SI4091 passwords after initial configuration and as regularly as required under your network security policies. For more information, see "Changing the System Passwords" on page 68.

Table 3. *Access Account Default Settings*

Account	Password	Description and Tasks Performed	Default
`user`	`user`	The User has no direct responsibility for SI4091 management. He or she can view all status information and statistics, but cannot make any configuration changes.	Disabled
`oper`	`oper`	The Operator manages all functions of the SI4091. The Operator can reset all ports except the management ports.	Disabled

Table 3. *Access Account Default Settings (continued)*

Account	Password	Description and Tasks Performed	Default
admin	admin	The super-user Administrator has complete access to all menus, information, and configuration commands on the SI4091, including the ability to change both the user and administrator passwords.	Enabled
USERID	PASSW0RD (with a zero)	An alternate administrator account. This admin-level account occupies end-user ID 1 (see "End-User Access Control" on page 75).	Enabled

For more information on defining or disabling switch access accounts, see "End-User Access Control" on page 75.

Administrator Password Recovery

Follow these steps to reset the password of the `admin` user to the default value:

Note: Password recovery process involves reloading the switch. Make sure to save any recent switch configuration changes before performing these steps.

1. Connect to the switch using the console port.
2. Reload the switch.
3. When the system displays Memory Test, press **<Shift + B>**. The Boot Management menu displays:

```
**** System Reset from boot iscli ****
Disable the Transceivers ...
Unmount the File System ...
Unmounting filesystem
Wait for umount to finish.Done
Waiting for I2C Transactions to Finish ...


U-Boot 2009.06 (Aug 21 2015 - 12:35:27) MPC83XX

Reset Status:

CPU: e300c4, MPC8378A, Rev: 2.1 at 792 MHz, CSB: 396 MHz
Board: Networking OS RackSwitch G8052
I2C: ready
DRAM: 1 GB

Memory Test ..........
```

4. Select **`C - Change configuration block`** from the Boot menu by entering **`C`**. When prompted for the configuration block, enter **`f`**:

```
Boot Menu Mode

Platform: Rack Switch G8052 (version 0.0.0.1)
FLASH: 256 MB
PCIE0: Link

Boot Management Menu
        I - Change booting image
        C - Change configuration block
        R - Boot in recovery mode (xmodem download of images to recover
switch)
        Q - Reboot
        E - Exit
Please choose your menu option: c

Currently using active configuration block
Enter configuration block: a, b or f (active, backup or factory): f
```

5. Enter **Q** to reboot the switch:

```
Boot Management Menu
        I - Change booting image
        C - Change configuration block
        R - Boot in recovery mode (xmodem download of images to recover
switch)
        Q - Reboot
        E - Exit
Please choose your menu option: q
Resetting the board.
```

6. After the reload is complete, log into the switch by using the default user `admin` with the default password `admin`.

7. Enter configuration mode (config). Copy the active configuration to the running configuration by using the **`copy active-config running-config admin-pw-bypass`** command.

```
SI 4091> ena
Enable privilege granted.
SI 4091# configure terminal
Enter configuration commands, one per line. End with Ctrl/Z.
SI 4091(config)# copy active-config running-config admin-pw-bypass
Loading to current configuration.
```

8. Use the **`show run`** command to confirm the configuration is recovered. Set the new admin and enable passwords. Save the running configuration to startup configuration.

```
SI 4091(config)# password
Changing admin password; validation required:
Enter current local admin password:
Enter new admin password (max 64 characters):
Re-enter new admin password:
New admin password accepted.

Password changed and applied, but not saved.
Notifying administrator to save changes.

SI 4091(config)# enable password ?
  WORD The UNENCRYPTED (cleartext) 'enable' password
SI 4091(config)# enable password admin1
SI 4091(config)# copy running-config startup-config
Confirm saving to FLASH (y/n) ? y
Copy running configuration to startup configuration
Switch is currently set to use factory default config block on next boot.
Do you want to change that to the active config block (y/n) ? y
Next boot will use active config block.
```

9. Make sure the boot configuration-block is active by using the **`show boot configuration-block`** command. If it is not active, change the boot configuration-block with the following command:

```
SI 4091(config)# boot configuration-block active
```

Configuring No-Prompt Mode

If you expect to administer the switch using Lenovo Switch Center or another browser-based interface, you need to turn off confirmation prompts. To accomplish this, use the command:

```
SI 4091(config)# [no] terminal dont-ask
```

In no-prompt mode, confirmation prompts are disabled for this and future sessions.

Chapter 3. Initial Setup

To help with the initial process of configuring your SI4091, the built-in software includes a Setup utility. The Setup utility prompts you step-by-step to enter all the necessary information for basic configuration of the device.

The SI4091 will automatically prompt you whether or not to run the Setup utility when factory default settings are detected. Setup can also be activated manually from the Command Line Interface (CLI).

Information Needed for Setup

Setup requests the following information:

- Basic system information such as date & time
- Optional configuration for internal port negotiation mode
- Optional configuration of IP parameters
  - IP address for each IP management interface
  - IP address for the default gateway

Default Setup Options

The Setup prompt appears automatically whenever you login as the system administrator under the factory default settings.

1. Connect to the SI4091 via one of the methods described in "Establishing a Connection" on page 27.

   After connecting, the login prompt will appear as shown here.

   ```
   Enter login username:
   ```

2. At the prompt, enter the administrator user name. If the system is set to factory defaults, type **USERID** and press <Enter>:

   ```
   Enter login username: USERID
   ```

3. When prompted, enter the administrator password. If the system is set to factory defaults, type **PASSW0RD** (with a zero) and press <Enter>:

   ```
   Enter login password: PASSW0RD
   ```

   Note: For security purposes, the text that you type for the password will not be displayed on the screen even though it is being processed.

4. After logging in with the factory default password, the system will prompt you to change the administrator password:

   ```
   Need to change default user password,
   Enter New Password (max 128 characters):
   ```

   Note: If this prompt does not appear, the password may have been changed from the factory default settings. If desired, return the SI4091 to its factory default configuration and start again.

5. Next, you will be prompted to verify the new administrator password.

   If you correctly enter the same new password, the administrator password will be changed and your login will continue. The system will then display a variety of system identification and operational information.

6. If factory default settings are in place (aside from the newly changed administrator password), the SI4091 will prompt whether or not you wish to run the Setup utility:

   ```
   The switch is booted with factory default configuration.
     To ease the configuration of the switch, a "Set Up" facility which
     will prompt you with those configuration items that are essential
     to the operation of the switch is provided.

   Would you like to run "Set Up" to configure the switch? [y/n]
   ```

7. Enter **y** to begin the initial configuration of the SI4091, or **n** to bypass the Setup utility.

   If you elect to bypass initial setup, you can manually restart the Setup utility at a later time (see "Stopping and Restarting Setup Manually" on page 46).

Stopping and Restarting Setup Manually

Follow these instructions to stop or restart the Setup utility manually.

Stopping Setup

To abort the Setup utility while it is running, press <Ctrl-C> at any Setup question. When you abort Setup, the system will prompt:

```
Would you like to run from top again? [y/n]
```

Enter **n** to abort Setup, or **y** to restart the Setup program at the beginning.

Restarting Setup

You can restart the Setup utility manually at any time by entering the following CLI commands:

```
SI 4091> enable
SI 4091# setup
```

The `enable` command initiates Privileged EXEC mode. The `setup` command can be executed in Privileged EXEC mode, and also in the Privileged EXEC configuration mode (at the `config` prompt).

Setup Part 1: Basic System Configuration

When Setup is started, the system prompts:

```
"Set Up" will walk you through the configuration of
  System Date and Time, Port Speed/Mode,
  VLANs, and IP interfaces. [type Ctrl-C to abort "Set Up"]
```

Note: The sample screens that appear in this document might differ slightly from the screens displayed by your system. Screen content varies based on the type of chassis unit that you are using and the firmware versions and options that are installed.

Next, the Setup utility prompts you to input basic system information.

1. Enter the year of the current date at the prompt:

   ```
   System Date:
   Enter year [2013]:
   ```

   Enter the four-digits that represent the year. To keep the current year, press <Enter>.

2. Enter the month of the current system date at the prompt:

   ```
   Enter month [11]:
   ```

   Enter the month as a number from 1 to 12. To keep the current month, press <Enter>.

3. Enter the day of the current date at the prompt:

   ```
   Enter day [15]:
   ```

   Enter the date as a number from 1 to 31. To keep the current day, press <Enter>.

   The system displays the date and time settings:

   ```
   System clock set to 18:55:36 Fri Nov 15, 2013.
   ```

4. Enter the hour of the current system time at the prompt:

   ```
   System Time:
   Enter hour in 24-hour format [18]:
   ```

   Enter the hour as a number from 00 to 23. To keep the current hour, press <Enter>.

5. Enter the minute of the current time at the prompt:

   ```
   Enter minutes [55]:
   ```

   Enter the minute as a number from 00 to 59. To keep the current minute, press <Enter>.

6. Enter the seconds of the current time at the prompt:

```
Enter seconds [37]:
```

Enter the seconds as a number from 00 to 59. To keep the current second, press <Enter>. The system then displays the date and time settings:

```
System clock set to 18:55:37 Fri Nov 15, 2013.
```

Setup Part 2: Port Configuration

Note: The sample screens that appear in this document might differ slightly from the screens displayed by your system. Screen content varies based on the type of chassis unit that you are using and the firmware versions and options that are installed.

1. Select the port to configure, or skip port configuration at the prompt:

```
Port Config:
Enter port (INT1-14, EXT1-10):
```

   If you wish to change settings for individual ports, enter the number of the port you wish to configure. To skip port configuration, press <Enter> without specifying any port.

2. Configure Gigabit Ethernet port flow parameters (if available).

   The system prompts:

```
Gig Link Configuration:
Port Flow Control:
```

3. Configure Gigabit Ethernet port autonegotiation mode (if available).

   If you selected a port that has a Gigabit Ethernet connector, the system prompts:

```
Port Auto Negotiation:
Current Port EXT1 autonegotiation:         on
Enter new value ["on"/"off"]:
```

   Enter **on** to enable port autonegotiation, **off** to disable it, or press <Enter> to keep the current setting.

   Notes:

   - In the current software, autonegotiation cannot be changed for external ports using the Setup utility under factory default conditions. Instead of a configuration prompt, an error message is displayed when external port configuration is attempted.
   - If you enable flow control in both directions and then disable autonegotiation, flow control will be disabled until you bring the link down and up.

4. The system prompts you to configure the next port:

```
Enter port (INT1-14, EXT1-10):
```

   When you are through configuring ports, press <Enter> without specifying any port. Otherwise, repeat the steps in this section.

Setup Part 3: IP Configuration

The system prompts for IP parameters.

IP Interfaces

The Setup utility allows one optional IPv4 management interface and/or one optional IPv6 management interface to be configured on the SI4091. The IP address assigned to each interface provides the SI4091 with an IP presence on their respective management networks. The administrator can use the configurated IP addresses to connect to the SI4091 for remote configuration.

Note: IP management interfaces operate under VLAN 4095, which is reserved for management functions and cannot be used in attached data domains.

1. Select the IP interface to configure, or skip interface configuration at the prompt:

```
IP Config:

IP interfaces:
Enter interface number: (125, 127)
```

If you wish to configure an IPv4 management interface, select interface 127. To configure an IPv6 management interface, select interface 125. To skip IP interface configuration, press <Enter> without typing an interface number and go to "Default Gateways" on page 52.

2. If configuring an IPv4 management interface (interface 127):

   a. Enter the IP address in IPv4 dotted decimal notation:

```
Enter new IP address:
```

   To keep the current setting (if any is displayed), press <Enter>.

   b. At the prompt, enter the IPv4 subnet mask in dotted decimal notation:

```
Enter new subnet mask:
```

   To keep the current setting (if any is displayed), press <Enter>.

   c. At the prompt, enter the VLAN for the interface:

```
Current VLAN: 4095
Enter new VLAN [1-4095]:
```

   To keep the current setting, press <Enter>.

   Note: In the current software, the IPv4 management interface VLAN cannot be changed using the Setup utility under factory default conditions. An error message is displayed when you attempt to change the setting.

   d. At the prompt, enter **y** to enable the IP interface, or **n** to leave it disabled:

```
Enable IP interface? [y/n]
```

3. If configuring an IPv6 management interface (interface 125):
   a. Enter the IP address in IPv6 hexidecimal notation (see "IPv6 Address Format" on page 175):

```
Enter new IP address:
```

      To keep the current setting (if any is displayed), press <Enter>.

   b. At the prompt, enter the IPv6 anycast address if desired:

```
Enter anycast if <IPv6 address> is anycast|<CR>:
```

      To keep the current setting (if any is displayed), press <Enter>.

   c. At the prompt, enter the IPv6 prefix length:

```
Enter new Prefix length [1-128]:
```

      To keep the current setting (if any is displayed), press <Enter>.

   d. At the prompt, enter the VLAN for the interface:

```
Current VLAN: 4095
Enter new VLAN [1-4095]:
```

      To keep the current setting, press <Enter>.

      Note: In the current software, the IPv4 management interface VLAN cannot be changed using the Setup utility under factory default conditions. An error message is displayed when you attempt to change the setting.

   e. At the prompt, enter **y** to enable the IP interface, or **n** to leave it disabled:

```
Enable IP interface? [y/n]
```

4. The system prompts you to configure another interface:

```
Enter interface number: (125, 127)
```

   Repeat the steps in this section until all desired management interfaces have been configured. When all interfaces have been configured, press <Enter> without specifying any interface number.

Default Gateways

1. At the prompt, select an IP default gateway for configuration, or skip default gateway configuration:

```
IP default gateways:
Enter default management gateway number: (3, 4)
```

   Enter the number for the IP default gateway to be configured. To skip default gateway configuration, press <Enter> without typing a gateway number and go to "Setup Part 4: Final Steps" on page 53.

   Note: In the current software, default gateway 4 cannot be configured using the Setup utility under factory default conditions. An error message is displayed when gateway 4 is selected.

2. At the prompt, enter the IPv4 address for the selected default gateway:

```
Current IP address:     0.0.0.0
Enter new IP address:
```

   Enter the IPv4 address in dotted decimal notation, or press <Enter> without specifying an address to accept the current setting.

3. At the prompt, enter **y** to enable the default gateway, or **n** to leave it disabled:

```
Enable default gateway? [y/n]
```

4. The system prompts you to configure another default gateway:

```
Enter default gateway number: (3, 4)
```

   Repeat the steps in this section until all default gateways have been configured. When all default gateways have been configured, press <Enter> without specifying any number.

Setup Part 4: Final Steps

1. When prompted, decide whether to restart Setup or continue to the final steps:

   ```
   Would you like to run from top again? [y/n]
   ```

   Enter **y** to restart the Setup utility from the beginning, or **n** to continue.

2. When prompted, decide whether you wish to review the configuration changes:

   ```
   Review the changes made? [y/n]
   ```

   Enter **y** to review the changes made during this session of the Setup utility. Enter **n** to continue without reviewing the changes. We recommend that you review the changes.

3. Next, decide whether to apply the changes at the prompt:

   ```
   Apply the changes? [y/n]
   ```

   Enter **y** to apply the changes, or **n** to continue without applying. Changes are normally applied.

4. At the prompt, decide whether to make the changes permanent:

   ```
   Save changes to flash? [y/n]
   ```

   Enter **y** to save the changes to flash. Enter **n** to continue without saving the changes. Changes are normally saved at this point.

5. If you do not apply or save the changes, the system prompts whether to abort them:

   ```
   Abort all changes? [y/n]
   ```

   Enter **y** to discard the changes. Enter **n** to return to the "`Apply the changes?`" prompt.

Optional Setup for Telnet Support

Note: This step is optional. Perform this procedure only if you are planning on connecting to the SI4091 through a remote Telnet connection.

Telnet is disabled by default. To change the setting, use the following configuration command:

```
SI 4091# configure terminal
SI 4091(config)# [no] access telnet enable
```

Chapter 4. Updating the System Software

The SI4091 software image is the executable code that directs system operation. A version of the image comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your SI4091.

The typical upgrade process for the software image consists of the following steps:

- Determine the version of software currently installed on your system.
- Get the latest version of software available for your system.
- Load a new software image and boot image onto an FTP, SFTP, or TFTP server on your network.
- Transfer the new images to your SI4091.
- Specify the new software image as the one which will be loaded into SI4091 memory the next time a system reset occurs.
- Reset the SI4091.

Detailed instructions for this typical upgrade process are covered in the rest of this chapter.

Determining the Current Software Version

To determine the software version currently used on the system, enter the following CLI command:

```
SI 4091# show boot
```

The software version will be shown in the resulting display.

Getting the Latest SI4091 Software

Note: This section includes references to IBM web sites and information in regard to obtaining service. IBM is Lenovo's preferred service provider for the System x, Flex System, and NeXtScale System products.

To get the latest version of software supported for your SI4091, go to the following website:

http://www.ibm.com/systems/support

Loading New Software to Your SI4091

CAUTION:
Although the standard software installation process as described in this section is all that is necessary in most cases, installing certain versions of ENOS requires additional, special steps to be taken prior to and/or after software installation. Check the *Release Notes* available for the specific version of the software you wish to install, and follow all applicable instructions. Failing to heed the full instructions in the *Release Notes* may cause unexpected behavior in the SI4091.

The SI4091 can store up to two different system software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

CAUTION:
When you upgrade the SI4091 software image, always load both the new boot image and the new software image before you reset the system. If you do not load a new boot image, your SI4091 might not boot properly (to recover, see "Recovering from a Failed Upgrade" on page 62).

To load a new software image to your SI4091, you will need the following:

- The image and boot software loaded on an FTP, SFTP, or TFTP file server on your network.

  Note: Be sure to download both the new boot file and the new image file.

- The hostname or IP address of the file server

  Note: The DNS parameters must be configured if specifying hostnames.

- The name of the new software image or boot file

When the software requirements are met, use the following procedure to download the new software to your SI4091:

1. Since FTP is disabled by default, you need to enable it. In Privileged EXEC mode, enter the following command:

```
SI 4091# access ftp enable
```

2. Copy the image to your switch, specifying the method for loading the software (FTP, SFTP, or TFTP) and the SI4091 destination (`image1`, `image2`, or `boot-image`) by entering the following command:

```
SI 4091# copy {tftp|ftp|sftp} {image1|image2|boot-image}
```

3. When prompted, enter the hostname or IP address of the file server.

```
Address or name of remote host: <name or IP address>
```

4. When prompted, enter the name of the new software file on the server.

```
Source file name: <filename>
```

   The exact form of the name will vary by server. However, the file location is normally relative to the FTP, SFTP, or TFTP directory (for example, `tftpboot`).

5. If required by the file server, enter the appropriate username and password.

6. The SI4091 will prompt you to confirm your request.

   Once confirmed, the software will begin loading into the SI4091.

7. When loading is complete, use the following commands to enter Global Configuration mode and select which software image (`image1` or `image2`) you want to run in system memory for the next reboot:

```
SI 4091# configure terminal
SI 4091(config)# boot image {image1|image2}
```

   The system will then verify which image is set to be loaded at the next reset:

```
Next boot will use switch software image1 instead of image2.
```

8. Reboot the SI4091 to run the new software:

```
SI 4091(config)# reload
```

   The system prompts you to confirm your request. Once confirmed, the system will reboot to use the new software.

Updating Software on vLAG Switches

When updating the software and boot images for switches configured with vLAG, first:

- Make sure that the spanning tree root switch is not one of the vLAG switches
- Shut down of ports should done under the port configuration
- Follow the shut down order of the ports

  a. ISL links

  b. vLAG links

  c. vLAG health check (MGT port)

Then follow this procedure to update the software on vLAG switches:

1. On Switch 2 (the original Secondary switch), shut down all links ISL, vLAG links, and vLAG HC. This is equivalent to powering off Switch 2.
   - All the traffic will failover to Switch 1 (the original Primary switch.).
   - After the shutdown of links on Switch 2, there will be N-S traffic loss of around ~0.16 seconds.
2. Upgrade Switch 2 with the new image. Use FTP, STFP, or TFTP to copy the new ENOS and boot images onto the switch. For more details, see "Loading New Software to Your SI4091" on page 58.
   - After Switch 2 comes up, vLAG HC will be up and vLAG mismatch will happen with vLAG ports down (since it is still Secondary).
   - The traffic will still be forwarding via Switch 1 (the original Primary switch).
3. On Switch 1 (the original Primary switch), shut down all links ISL, vLAG links, and vLAG HC. This is equivalent to powering off Switch 1 (the original Primary switch)
   - All the traffic will failover to Switch 2, which will assume the vLAG operation role of Primary.
   - After all the links are up on Switch 2, there will be N-S traffic loss of around ~70 seconds due to spanning trees reconverging.
4. Upgrade Switch 1 (the original Primary switch with the new ENOS image. Use FTP, STFP, or TFTP to copy the new ENOS and boot images onto the switch. For more details, see "Loading New Software to Your SI4091" on page 58.
   - After Switch 1 comes up, vLAG HC, ISL, and vLAG links will be up, and Switch 1 will assume the vLAG operation role of Secondary.
   - All the traffic will now follow the hash and load balance settings between Switch 1 and Switch 2.
   - There will be N-S traffic loss of around ~0.05 seconds.
5. Change the operational role of the vLAG switches to match the final topology by reloading Switch 2.
   - There will be N-S traffic loss of around ~0.102 seconds.

- Switch 1 will reassume the vLAG Primary role and Switch 2 will reassume the vLAG Secondary role.

6. Make sure that Switch 1 is now the vLAG primary switch and Switch 2 is now the vLAG secondary switch using the following command:

```
SI 4091> show vlag information
```

Recovering from a Failed Upgrade

The Boot Management menu allows you to perform fundamental device management operations, such as selecting which software image will be loaded, resetting the SI4091 to factory defaults, or recovering from a failed software download.

Use the following procedure to recover from a failed software upgrade.

1. Connect a PC to the serial Console port of the SI4091.
2. Open a terminal emulator program that supports XModem Download (for example, HyperTerminal, SecureCRT, PuTTY) and select the following port characteristics:
   - Speed: 9600 bps
   - Data Bits: 8
   - Stop Bits: 1
   - Parity: None
   - Flow Control: None
3. To access the Boot Management menu, you must interrupt the boot process from the Console port. Boot the SI4091, and when the system begins displaying `Memory Test` progress (a series of dots), press **<Shift B>**.

   The Boot Management menu will appear:

```
Resetting the System ...
Memory Test ..............................

Boot Management Menu
    1 - Change booting image
    2 - Change configuration block
    3 - Boot in recovery mode (tftp and xmodem download of images to
        recover switch)
    5 - Reboot
    6 - Exit
```

4. Select **3** to boot in recovery mode. The following menu will appear:

```
Please select one of the following options:
    T) Configure networking and tftp download an image
    X) Use xmodem 1K to serial download an image
    P) Physical presence (low security mode)
    R) Reboot
    E) Exit
```

5. Select **T** to configure networking options for TFTP image download.

   When prompted, specify which management port type (internal or external) is to be used, the switch management IP address, netmask, and gateway, and the FTPT server IP address.

6. Select **X** to begin xmodem serial download of the switch image.

7. When you see the following message, change the port characteristics in your terminal emulation software to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```

8. Press <**Enter>** to set the system into download accept mode. When the readiness meter displays (a series of "C" characters), start XModem on your terminal emulator.

9. Select the Boot Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```
yzModem - CRC mode, 62494(SOH)/0(STX)/0(CAN) packets, 6 retries

Extracting images ... Do *NOT* power cycle the switch.

**** VMLINUX ****

Un-Protected 10 sectors

Erasing Flash.............. done

Writing to Flash.............done

Protected 10 sectors

**** RAMDISK ****

Un-Protected 44 sectors

Erasing Flash.............................................. done

Writing to Flash.............................................done

Protected 44 sectors

**** BOOT CODE ****

Un-Protected 8 sectors

Erasing Flash........... done

Writing to Flash...........done

Protected 8 sectors
```

10. When you see the following message, change the port characteristics in your terminal emulation software to 9600 bps:

```
## Switch baudrate to 9600 bps and press ESC ...
```

11. Press the Escape key (<**Esc>**) to re-display the Boot Management menu.

12. Select **3** to start a new `XModem Download`. When you see the following message, change the port characteristics in your terminal emulation software to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```

13. Press <**Enter>** to continue the download.
14. Select the ENOS Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```
yzModem - CRC mode, 27186(SOH)/0(STX)/0(CAN) packets, 6 retries

Extracting images ... Do *NOT* power cycle the switch.

**** Switch OS ****


Please choose the Switch OS Image to upgrade [1|2|n] :
```

15. Select the image number to load the new image (1 or 2). It is recommended that you select 1. A message similar to the following is displayed:

```
Switch OS Image 1 ...

Un-Protected 27 sectors

Erasing Flash.............................. done

Writing to Flash.............................done

Protected 27 sectors
```

16. When you see the following message, change the port characteristics in your terminal emulation software to 9600 bps:

```
## Switch baudrate to 9600 bps and press ESC ...
```

17. Press the Escape key (<**Esc>**) to re-display the Boot Management menu.
18. Select **6** to exit and boot using the new image.

Part 2: Securing the SI4091

Chapter 5. Administrative Security

This chapter discusses different methods of securing local and remote administration on the SI4091 System Interconnect Module (SI4091):

Changing the System Passwords

Access to the SI4091 command line interface is controlled through the use of a login. Once you are connected to the SI4091, you are prompted to enter a login name and its associated password. By default, two logins are available:

- The administrator

  The administrator login has complete access to all menus, information, and configuration commands, including the ability to change both the user and administrator passwords.

  The default username is: `admin`

  The default password is: `admin`

- An alternate administrator:

  The default username is: `USERID`

  The default password is: `PASSW0RD` (with a zero)

  This admin-level account occupies end-user ID 1 (see "End-User Access Control" on page 75).

It is recommended that you change the passwords after initial setup and as regularly as required under your network security policies. The administrator account can change all passwords, as shown in the following procedure.

1. Log-in to the SI4091.

   When prompted, enter the administrator username and password:

2. Access the configuration mode:

```
SI 4091> enable
SI 4091# configure terminal
SI 4091(config)#
```

   The `enable` command initiates Privileged EXEC mode, and the `configure terminal` command readies the SI4091 for basic configuration. The system prompt changes to indicate the current command mode.

3. Set the new administrator "`admin`" password:

```
SI 4091(config)# access user administrator-password
```

   This command will prompt for required information: current administrator password, new administrator password (up to 128 characters), and confirmation of the new administrator password.

4. Set the new alternate administrator "`USERID`" password:

```
SI 4091(config)# access user 1 password
```

   This command will prompt for required information: current administrator password, new administrator password (up to 128 characters), and confirmation of the new administrator password.

Secure Shell and Secure Copy

Because using Telnet does not provide a secure connection for managing a SI4091, Secure Shell (SSH) and Secure Copy (SCP) features have been included for SI4091 management. SSH and SCP use secure tunnels to encrypt and secure messages between a remote administrator and the SI4091.

- **SSH** is a protocol that enables remote administrators to log securely into the SI4091 over a network to execute management commands. The SI4091 supports SSH version 2.0 and above.

  By default, SSH is *enabled* on the SI4091.

- **SCP** is typically used to copy files securely from one device to another. SCP uses SSH for encryption of data on the network. On the SI4091, SCP is used to download and upload the system configuration via secure channels.

  By default, SCP is *disabled* on the SI4091.

Configuring SSH/SCP

By default, SSH is enabled and SCP is disabled. To change the setting, use the following procedures.

Note: SCP requires SSH to remain enabled.

To Enable or Disable the SSH Feature

Begin a Telnet session from the Console port and enter the following commands:

```
SI 4091> enable
SI 4091# configure terminal
SI 4091(config)# [no] ssh enable
```

The `enable` command initiates Privileged EXEC mode, and the `configure terminal` command readies the SI4091 for basic configuration. The system prompt changes to indicate the current command mode.

To Enable or Disable SCP Apply and Save

Enter the following command from the configuration prompt to enable or disable the SCP `putcfg_apply` and `putcfg_apply_save` commands:

```
SI 4091(config)# [no] ssh scp-enable
```

Configuring the SCP Administrator Password

To configure the SCP-only administrator password, enter the following command from the configuration prompt (the default SCP-only password is `admin`):

```
SI 4091(config)# [no] ssh scp-password
Changing SCP-only Administrator password; validation required...
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```

Using SSH and SCP Client Commands

This section shows the format for using some common client commands.

To Log In to the SI4091 from the Client

Syntax:

```
>> ssh [-4|-6] <SI4091 IP address>
   -or-
>> ssh [-4|-6] <login name>@<SI4091 IP address>
```

Note: The `-4` option (the default) specifies that an IPv4 address on the SI4091 will be used. The `-6` option specifies IPv6.

Example**:**

```
>> ssh scpadmin@205.178.15.157
```

To Copy the SI4091 System Configuration File to the SCP Host

Syntax**:**

```
>> scp [-4|-6] <username>@<SI4091 IP address>:getcfg <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getcfg ad4.cfg
```

To Load a SI4091 System Configuration File from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<SI4091 IP address>:putcfg
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg
```

To Save the Configuration

When loading a configuration file to the SI4091, the commands in the file are automatically applied to the current configuration. However, the commands are not saved, and will be lost when the SI4091 is rebooted unless saved to the system's FLASH memory. The appropriate commands may be entered using SCP.

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<SI4091 IP address>:putcfg_apply_save
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg_apply_save
```

The `putcfg_apply_save` command is usually preferred over `putcfg` because the SCP session is not in an interactive mode.

To Copy the SI4091 Software Image and Boot Files to the SCP Host

Syntax**:**

```
>> scp [-4|-6] <username>@<SI4091 IP address>:getimg1 <local filename>
>> scp [-4|-6] <username>@<SI4091 IP address>:getimg2 <local filename>
>> scp [-4|-6] <username>@<SI4091 IP address>:getboot <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getimg1 6.1.0_os.img
```

To Load SI4091 System Configuration Files from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<SI4091 IP address>:putimg1
>> scp [-4|-6] <local filename> <username>@<SI4091 IP address>:putimg2
>> scp [-4|-6] <local filename> <username>@<SI4091 IP address>:putboot
```

Example:

```
>> scp 6.1.0_os.img scpadmin@205.178.15.157:putimg1
```

SSH and SCP Encryption of Management Messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication: 1024-bit RSA host key
- Key Exchange: diffie-hellman-group1-sha1, diffie-hellman-group14-sha1
- Encryption: 3des-cbc, aes128-cbc, aes128-ctr, arcfour, arcfour128, arcfour256, blowfish-cbc, rijndael128-cbc
- MAC: hmac-md5, hmac-md5-96, hmac-sha1, hmac-sha1-96
- User Authentication: Local password authentication, LDAP, RADIUS, TACACS+

Generating the RSA Host Key for SSH Access

To support the SSH feature, an RSA host key is required. The host key is 1024 bits and is used to identify the SI4091.

When SSH is first enabled, the SI4091 automatically generates the RSA host key and stores it in FLASH memory.

To manually create an RSA host key, enter the following command via the SI4091 Console port (the command is not available via external Telnet connection)

```
SI 4091(config)# ssh generate-host-key
```

When the SI4091 reboots, it will retrieve the host key from the FLASH memory.

Notes:

- The switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is performing key generation at that time. Also, key generation will fail if an SSH/SCP client is logging in at that time.
- Because the switch software only generates RSA keys, if there is already a DSA-based SSH key on the switch, this key will remain on the switch and not be replaced until you run the `ssh generate-host key` command to generate an RSA key.

SSH/SCP Integration with RADIUS Authentication

SSH/SCP is integrated with RADIUS authentication. After the RADIUS server is enabled on the SI4091, all subsequent SSH authentication requests will be redirected to the specified RADIUS servers for authentication. The redirection is transparent to the SSH clients.

SSH/SCP Integration with TACACS+ Authentication

SSH/SCP is integrated with TACACS+ authentication. After the TACACS+ server is enabled on the SI4091, all subsequent SSH authentication requests will be redirected to the specified TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

Secure FTP

The SI4091 supports Secure FTP (SFTP). SFTP uses Secure Shell (SSH) to transfer files. SFTP encrypts both commands and data, and prevents passwords and sensitive information from being transmitted openly over the network. The following command is an example of SFTP support on the SI4091:

```
SI 4091# copy sftp {image1|image2|boot-image} [mgt-port|extm-port]
```

This command copies a software file from the SFTP server to the SI4091.

End-User Access Control

Enterprise NOS allows an administrator to define end-user accounts that permit end users to perform operation tasks via the SI4091 CLI commands. Once end-user accounts are configured and enabled, the SI4091 requires username/password authentication.

For example, an administrator can assign a user, who can then log into the SI4091 and perform operational commands (effective only until the next SI4091 reboot).

Considerations for Configuring End User Accounts

- A maximum of 20 user IDs are supported on the SI4091.
- Enterprise NOS supports end-user account for the Console port, Telnet, and SSH version 2 (or above) access to the SI4091.
- If RADIUS authentication is used, the user password on the Radius server will override the user password on the SI4091. Also note that the password change command modifies only the user password on the SI4091 and has no effect on the user password on the Radius server. Radius authentication and user password cannot be used concurrently to access the SI4091.
- Passwords can be up to 64 characters in length for Telnet, SSH, Console, and Web access.

Strong Passwords

The administrator can require use of Strong Passwords for users to access the SI4091. Strong Passwords enhance security because they make password guessing more difficult.

The following rules apply when Strong Passwords are enabled:

- Minimum length: 8 characters; maximum length: 14 characters
- Must contain at least one uppercase alphabet
- Must contain at least one lowercase alphabet
- Must contain at least one number
- Must contain at least one special character:

  Supported special characters: ! “ # % & ‘ ( ) ; < = >> ? [\ ] * + , - . / : ^ _ { | } ~

When strong password is enabled, users can still access the SI4091 using the old password but will be advised to change to a strong password while attempting to log in.

Strong password requirement can be enabled using the following command:

```
SI 4091 # access user strong-password enable
```

The administrator can choose the number of days allowed before each password expires. When a strong password expires, the user is allowed to log in one last time (last time) to change the password. A warning provides advance notice for users to change the password.

User Access Control Commands

The end-user access control commands allow you to configure end-user accounts.

Up to 20 user IDs can be configured. Use the following commands to define any user name and set the user password at the resulting prompts:

```
SI 4091(config)# access user <user number (1-20)> name <1-64 characters>
SI 4091(config)# access user <user number (1-20)> password

Changing user1 password; validation required:
Enter current admin password: <current administrator password>
Enter new user1 password: <new user password of up to 128 characters>
Re-enter new user1 password: <new user password>
New user1 password accepted.
```

Note: User 1 is pre-configured for administrator-level access, with a login name of USERID and password of PASSW0RD (with a zero). This user ID cannot be disabled, nor can the name or access level be changed. Only the password can be changed.

Defining a User's Access Level

The end user is by default assigned to the user access level (also known as class of service, or CoS). CoS for all user accounts have global access to all resources except for User CoS, which has access to view only resources that the user owns. For more information, see Table 7 on page 92.

To change the user's level, select one of the following options:

```
SI 4091(config)# access user <user number (1-20)> level
{user|operator|administrator}
```

Note: The administrator access level for User 1 cannot be changed.

Validating a User's Configuration

```
SI 4091# show access user uid <user number (1-20)>
```

Enabling or Disabling a User

An end user account must be enabled before the SI4091 recognizes and permits login under the account. Once enabled, the SI4091 requires any user to enter both username and password.

```
SI 4091(config)# [no] access user <user number (1-20)> enable
```

Note: User 1 cannot be disabled.

Listing Current Users

The following command displays defined user accounts and whether or not each user is currently logged into the SI4091.

```
SI 4091# show access user

Usernames:
  user     - disabled   - offline
  oper     - disabled   - offline
  admin    - Always Enabled - online     1 session.
Current User ID table:
  1: name USERID  , ena, cos admin   , password valid, offline
```

Logging In to an End User Account

Once an end user account is configured and enabled, the user can login to the SI4091 by using the username/password combination. The degree of SI4091 access is determined by the Class of Service established for the end user account.

Boot Strict Mode

The implementations specified in this section are compliant with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131A.

The SI4091 can operate in two boot modes:

- Compatibility mode (default): This is the default boot mode. This mode may use algorithms and key lengths that are not allowed or acceptable by the NIST SP 800-131A specification. This mode is useful in maintaining compatibility with previous releases and in environments that have lesser data security requirements.
- Strict mode: Encryption algorithms, protocols, and key lengths in strict mode are compliant with the NIST SP 800-131A specification.

When in boot strict mode, the SI4091 uses Secure Sockets Layer (SSL)/Transport Layer Security (TLS) 1.2 protocols to ensure confidentiality of the data to and from the SI4091.

By default, Telnet, and SNMPv1 and SNMPv2 are disabled on the SI4091. In strict mode, if the security policy on the SI4091 is set to `secure`, a warning message will appear in the system log if one or more of these protocols are enabled.

Before enabling strict mode, ensure the following:

- Ensure that NIST Strict compliance is enabled on the Chassis Management Module (CMM)
- All connected SI4091 devices and Lenovo switches must run Enterprise NOS 8.4.
- The supported protocol versions and cryptographic cipher suites between clients and servers must be compatible. For example: if using SSH to connect to the SI4091, ensure that the SSH client supports SSHv2 and a strong cipher suite that is compliant with the NIST standard.
- It is recommended that protocols which are not NIST SP 800-131A compliant be disabled or remain unused.
- Only SSHv2 or higher may be used.
- Save the current configuration, if any, in a location external to the SI4091. When the SI4091 reboots, both the startup and running configuration are lost.
- It is recommended that only protocols or algorithms which are compliant with the NIST SP 800-131A specification be used or enabled on the SI4091. Please see the NIST SP 800-131A publication for details. Table 4 on page 79 lists the acceptable protocols and algorithms.

Table 4. *Acceptable Protocols and Algorithms*

Protocol and Function	Strict Mode Algorithm	Compatibility Mode Algorithm
Certificate Generation	RSA-2048 SHA-256	RSA 2048 SHA 256
Certificate Acceptance	RSA 2048 or higher SHA 224 or higher	RSA SHA, SHA2
LDAP	LDAP does not comply with NIST SP 800-131A specification. When in strict mode, LDAP is disabled. However, it can be enabled, if required.	Acceptable
RADIUS	RADIUS does not comply with NIST SP 800-131A specification. When in strict mode, RADIUS is disabled. However, it can be enabled, if required.	Acceptable
Random Number Generator	NIST SP 800-90A AES CTR DRBG	NIST SP 800-90A AES CTR DRBG
Secure NTP	Secure NTP does not comply with NIST SP 800-131A specification. When in strict mode, secure NTP is disabled. However, it can be enabled, if required.	Acceptable
SLP	SHA-256 or higher RSA/DSA 2048 or higher	
SNMP	SNMPv3 only AES-128-CFB-128/SHA1 **Note:** Following algorithms are acceptable if you choose to support old SNMPv3 factory default users: AES-128-CFB/SHA1 DES/MD5 AES-128-CFB-128/SHA1	SNMPv1, SNMPv2, SNMPv3 DES/MD5, AES-128-CFB-128/SHA1
SSH/SFTP		
Host Key	SSH-RSA	SSH-RSA
Key Exchange	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 RSA2048-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA1	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 ECDH-SHA2-NISTP192 RSA2048-SHA256 RSA1024-SHA1 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA1 DIFFIE-HELLMAN-GROUP14-SHA1 DIFFIE-HELLMAN-GROUP1-SHA1

Table 4. *Acceptable Protocols and Algorithms (continued)*

Protocol and Function	Strict Mode Algorithm	Compatibility Mode Algorithm
Encryption	AES128-CTR AES128-CBC 3DES-CBC	AES128-CTR AES128-CBC RIJNDAEL128-CBC BLOWFISH-CBC 3DES-CBC ARCFOUR256 ARCFOUR128 ARCFOUR
MAC	HMAC-SHA1 HMAC-SHA1-96	HMAC-SHA1 HMAC-SHA1-96 HMAC-MD5 HMAC-MD5-96
TACACS+	TACACS+ does not comply with NIST SP 800-131A specification. When in strict mode, TACACS+ is disabled. However, it can be enabled, if required.	Acceptable

Acceptable Cipher Suites

The following cipher suites are acceptable (listed in the order of preference) when the SI4091 is in compatibility mode:

Table 5. *List of Acceptable Cipher Suites in Compatibility Mode*

Cipher ID	Key Exchange	Auth.	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0xC011	ECDHE	RSA	RC4	SHA1	SSL_ECDHE_RSA_WITH_RC4_128_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x0005	RSA	RSA	RC4	SHA1	SSL_RSA_WITH_RC4_128_SHA
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES-128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA

The following cipher suites are acceptable (listed in the order of preference) when the SI4091 is in strict mode:

Table 6. *List of Acceptable Cipher Suites in Strict Mode*

Cipher ID	Key Exchange	Auth.	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES-128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA

Configuring Strict Mode

To change the SI4091 mode to or from boot strict mode, use the following Privileged EXEC Configuration mode command:

```
SI 4091(config)# [no] boot strict enable
```

When strict mode is enabled, you will see the following message:

```
Warning, security strict mode limits the cryptographic algorithms used by
secure protocols on this switch. Please see the documentation for full
details, and verify that peer devices support acceptable algorithms
before enabling this mode. The mode change will take effect after
reloading the switch and the configuration will be wiped during the
reload. System will enter security strict mode with default factory
configuration at next boot up.

Do you want SNMPV3 support old default users in strict mode (y/n)?
```

For SNMPv3 default users, see "SNMP Version 3" on page 259.

When strict mode is disabled, the following message is displayed:

```
Warning, disabling security strict mode. The mode change will take effect
after reloading the switch.
```

You must reboot the SI4091 for the boot strict mode enable/disable to take effect.

Boot Strict Mode Limitations

Consider the following limitation and restrictions if you need to operate the SI4091 in boot strict mode:

- Power ITEs and High-Availability features do not comply with NIST SP 800-131A specification.
- The SI4091 will not discover Platform agents/Common agents that are not in strict mode.
- Limited functions of the SI4091 managing Windows will be available.

Protected Mode

Protected Mode settings allow the SI4091 administrator to block the management module from making configuration changes that affect system operation. The SI4091 retains control over those functions.

The following management module functions are disabled when Protected Mode is turned on:

- External Ports: Enabled/Disabled
- External management over all ports: Enabled/Disabled
- Restore Factory Defaults
- New Static IP Configuration

In this release, configuration of the functions listed above are restricted to the local SI4091 when you turn Protected Mode on. In future releases, individual control over each function may be added.

Note: Before you turn Protected Mode on, make sure that external management (Telnet) access to one of the IP interfaces on the SI4091 is enabled.

Use the following Privileged EXEC Configuration mode command to turn Protected Mode on:

```
SI 4091(config)# protected-mode enable
```

If you lose access to the SI4091 through the external ports, use the Console port to connect directly to the SI4091, and configure an IP interface with Telnet access to reestablish access.

Generating a Certificate Signing Request

Before a digital certificate can be signed by a Certificate Authority (CA), it needs to be created. The generation of a certificate involves creating a Certificate Signing Request (CSR). The CSR includes various information related to the device and a public key. The public key is included in the CSR file itself and the private key associated with the public key is generated separately and kept private. The CSR can then be exported to a remote device to be signed by a CA.

1. Create an HTTPS CSR defining the information you want to be used in the various fields:

```
SI 4091(config)# access https generate-csr
Country Name (2 letter code) []:  <country code>
State or Province Name (full name) []:  <state>
Locality Name (eg, city) []:  <city>
Organization Name (eg, company) []:  <company>
Organizational Unit Name (eg, section) []:  <org. unit>
Common Name (eg, YOUR name) []:  <name>
Email (eg, email address) []:  <email address>
Confirm Generate CSR? [y/n]: y
.........................................+++
.............+++
Cert Req generated successfully
```

2. To verify the CSR you can use the following command:

show https host-csr [pem-format|txt-format]

```
SI 4091> show https host-csr txt-format

Certificate Request:
    Data:
        Version: 0 (0x0)
        Subject: C=US, ST=Cali, L=Santa Barbara, O=Lenovo, OU=Sales, CN=www.zagat.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
            RSA Public Key: (2048 bit)
                Modulus (2048 bit):
                    00:b5:05:f6:d5:ad:ab:f2:1d:a9:57:c4:bc:84:1b:
                    c6:bc:cd:04:95:ea:ad:ec:4a:44:3a:6e:42:9f:39:
                    96:14:11:a7:8e:3e:6f:da:9a:42:c6:c4:62:a1:33:
                    0e:a8:d3:6a:21:ce:f3:3c:4f:c1:8d:d1:e7:9e:c7:
                    29:04:ea:c6:7d:54:9a:4e:10:24:10:38:45:c6:4b:
                    13:19:f2:dd:8a:83:3f:5c:cf:8b:85:a7:2a:b0:eb:
                    7a:26:1f:4c:94:47:01:81:6a:59:d5:f5:d6:7e:3b:
                    b5:bc:e4:3f:6d:dd:84:15:07:61:93:e0:d1:40:f8:
                    9d:15:d0:a6:e1:9b:a4:ab:85:b5:2b:f0:56:e9:ef:
                    36:43:2b:aa:be:1b:63:3c:fd:74:ab:78:76:53:12:
                    e6:65:4c:0d:07:91:df:b3:91:96:f4:55:f7:37:73:
                    8c:f6:77:d7:9d:2b:a5:bd:17:3f:11:f2:85:4b:d6:
                    b4:1d:3f:70:1f:13:bb:5e:2e:4c:a8:ad:6a:7f:11:
                    36:97:a6:25:0a:87:66:31:c9:92:59:03:31:5d:ff:
                    df:c6:aa:93:7c:51:9f:8e:1b:6f:2a:be:c4:4c:66:
                    d6:2c:4b:6d:e6:ae:4e:02:82:fc:fa:a1:de:3b:c9:
                    24:25:d5:6e:15:15:18:ce:9b:a6:98:ad:0c:32:1f:
                    94:01
                Exponent: 65537 (0x10001)
        Attributes:
            a0:00
    Signature Algorithm: sha256WithRSAEncryption
        24:26:dd:96:49:47:9d:78:74:48:9b:63:4c:32:f0:78:da:7d:
        82:c9:17:6d:7e:93:38:60:94:d5:02:c1:31:dc:42:69:f5:57:
        46:a8:44:5a:99:ea:55:d3:99:bf:f0:48:3b:ef:60:fd:50:e6:
        33:cd:89:86:d3:51:97:f2:d1:68:6f:88:8c:e7:0f:3e:19:2a:
        f4:ea:6b:dc:05:24:d7:98:cd:a3:d3:c3:ef:03:93:8b:3f:fe:
        75:5e:67:f1:48:b6:20:a6:ff:ae:5a:25:41:7f:e4:c8:48:d4:
        63:37:16:98:9e:2d:1b:b6:65:7a:0d:90:87:07:19:f0:02:17:
        3a:3e:fd:f0:40:3e:a4:0f:53:97:9b:d5:18:22:78:f3:07:94:
        63:be:f9:f2:5c:23:6d:0f:22:d1:17:db:38:24:5c:6b:7b:e0:
        41:a6:51:28:30:2c:f4:1d:62:6c:06:f2:4c:0c:5b:79:51:13:
        73:f8:88:ba:2e:05:98:5d:41:5e:9d:58:b1:0c:8f:fc:f2:79:
        d5:30:7c:95:e9:ff:9a:cc:dd:d9:4c:2e:98:32:5a:ab:cd:59:
        a4:37:a5:38:03:4e:e7:27:dc:14:c8:75:9d:ca:e0:62:37:02:
        19:17:16:e3:92:c0:c3:16:13:26:c9:40:d7:ec:f2:8c:8e:fc:
        1a:dc:27:4c
```

```
SI 4091> show https host-csr pem-format

-----BEGIN CERTIFICATE REQUEST-----
MIICtDCCAZwCAQAwbzELMAkGA1UEBhMCVVMxEzARBgNVBAgMCkNhbGlmb3JuaWEx
ETAPBgNVBAcMCFNhbiBKb3NlMQwwCgYDVQQKDANBQkMxFDASBgNVBAsMC0VuZ2lu
ZWVyaW5nMRQwEgYDVQQDDAt3d3cuYWJjLmNvbTCCASIwDQYJKoZIhvcNAQEBBQAD
ggEPADCCAQoCggEBAMEnVJBSnIY9OGiCcH+xmYKpWga7E5j9JSK9JU57Md7NofJ2
FvQ8hfPO8b4bzLQzKbNBxGc59BJjZJ5w8eGKRDCjlIf1uIAgg3Gs8ZK1FozOUJZN
xbtYBx6QrTBYmXdHStQ7CQ9sfWhnEnusnvc8bxNlukyuEcFsAUdz93r1sEfN3cDe
/bO43l7GmvhTEdmfFvAfgi9b9RDqUjliATpgS3+a2kwhjvHCTeveQN1/MYQZvbJo
V4qq+pgQOt9ZJOMDrGQ0YmO1p84+GdxXVwGePCOvCRLESsq5rQb3zPSVvWnTsq0G
gURvbV+VQN9dI9lANZGZJi6BRNIRdBen/dH0KRcCAwEAAaAAMA0GCSqGSIb3DQEB
BQUAA4IBAQCSLDOrOnl7kaZri2OjDpzgiiG+9Skde3MehaklddfZnCkT1ALL3ZXY
xWwYnvF5jAgnHhxRJbPOzwHNDWMtZiiNOTHyzHVptsyRBv70Kb8odJmuyKWDqunJ
Ho1hHe63a7MRLFkQ+6io3kGrmq1bdM5U6xvvS+0ZXXUaiK1p/lNLOrsYk45D01Az
YHhcdRQtFUbQxqbirpi0jLsi82X7JCNQ2XCP6dhphkWKI6wsCvmlJdazW+V/gH/X
wqMkNF8mkodzlhC+1COd2yzSXxqpG/Xf0TRF9SAyN5vK4NDDzZu+iPvh6RkXXeNV
neyr2J5JENyGORPynuVoHWuzEy+5GUHa
-----END CERTIFICATE REQUEST-----
```

3. Export the CSR file to an external server:

```
SI 4091(config)# copy cert-request tftp

Port type ["DATA"/"MGT"/"EXTM"]: <port type>
Address or name of remote host: <hostname or IPv4 address>
Destination file name: <path and filename on the remote server>

Certificate request successfully tftp'd to...
```

Maintenance Mode

There are times when Lenovo support needs to access your switch in maintenance mode. To enable this, enter the command:

```
SI 4091(config)# maint-internal
```

When prompted, enter the admin password.

The Lenovo support person will then enter the maintenance mode password.

This introduces a second level of administration authorization before the Lenovo support representative enters the maintenance mode password, making the switch more secure and available for enhanced debugging.

Chapter 6. Authentication & Authorization Protocols

Secure SI4091 management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured IPv4 management and device access:

- "RADIUS Authentication and Authorization" on page 90
- "TACACS+ Authentication" on page 94
- "LDAP Authentication and Authorization" on page 99

Note: Enterprise NOS 8.4 does not support IPv6 for RADIUS, TACACS+ or LDAP.

RADIUS Authentication and Authorization

The SI4091 supports the RADIUS (Remote Authentication Dial-in User Service) method to authenticate and authorize remote administrators for managing the SI4091. This method is based on a client/server model. The Remote Access Server (RAS)—the SI4091—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes UDP over IP (based on RFC 2138 and 2866)
- A centralized server that stores all the user authorization information
- A client (in this case, the SI4091 is the client)

The SI4091—acting as the RADIUS client—communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the SI4091) and the back-end RADIUS server.

How RADIUS Authentication Works

1. Remote administrator connects to the SI4091 and provides user name and password.
2. Using Authentication/Authorization protocol, the SI4091 sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the SI4091 to grant or deny administrative access.

Configuring RADIUS on the SI4091

Use the following procedure to configure Radius authentication on your SI4091.

1. Configure the IPv4 addresses of the Primary and Secondary RADIUS servers, and enable RADIUS authentication.

```
SI 4091(config)# radius-server primary-host 10.10.1.1
SI 4091(config)# radius-server secondary-host 10.10.1.2
SI 4091(config)# radius-server enable
```

2. Configure the RADIUS secret.

```
SI 4091(config)# radius-server primary-host 10.10.1.1 key
     <1-32 character secret>
SI 4091(config)# radius-server secondary-host 10.10.1.2 key
<1-32 character secret>
```

CAUTION:
If you configure the RADIUS secret using any method other than through the Console port, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default UDP port number used to listen to RADIUS.

   The well-known port for RADIUS is 1645.

```
SI 4091(config)# radius-server port <UDP port number>
```

4. Configure the number retry attempts for contacting the RADIUS server, and the timeout period.

```
SI 4091(config)# radius-server retransmit 3
SI 4091(config)# radius-server timeout 5
```

RADIUS Authentication Features in Enterprise NOS

Enterprise NOS supports the following RADIUS authentication features:

- Supports RADIUS client on the SI4091, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows a RADIUS secret password of up to 32 characters.
- Supports *secondary authentication server* so that when the primary authentication server is unreachable, the SI4091 can send client authentication requests to the secondary authentication server. Use the following command to show the currently active RADIUS authentication server:

```
SI 4091# show radius-server
```

- Supports user-configurable RADIUS server retry and time-out values:
  - Time-out value = 1-10 seconds
  - Retries = 1-3

  The SI4091 will time out if it does not receive a response from the RADIUS server within 1-10 seconds. The SI4091 automatically retries connecting to the RADIUS server 1-3 times before it declares the server down.
- Supports user-configurable RADIUS application port. The default is UDP port 1645. UDP port 1812, based on RFC 2138, is also supported.
- Allows network administrator to define privileges for one or more specific users to access the SI4091 at the RADIUS user database.

User Accounts

The user accounts listed in Table 7 can be defined in the RADIUS server dictionary file.

Table 7. *User Access Levels*

Account	Description and Tasks Performed	Default Password
Level: User Name: `user`	The User has no direct responsibility for SI4091 management. They can view all system status information and statistics but cannot make any configuration changes to the system.	(disabled)
Level: Operator Name: `oper`	In addition to User capabilities, the Operator has limited system management access, including the ability to make temporary, operational configuration changes to some SI4091 features, and to reset SI4091 ports (other than management ports).	(disabled)
Level: Admin Name: `admin`	The super-user Administrator has complete access to all menus, information, and configuration commands on the SI4091, including the ability to change both the user and administrator passwords.	`admin`
Level: Admin Name: `USERID`	An alternate administrator account. This admin-level account occupies end-user ID 1 (see "End-User Access Control" on page 75).	`PASSW0RD` (with a zero)

RADIUS Attributes for Enterprise NOS User Privileges

When the user logs in, the SI4091 authenticates his/her level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the remote user is successfully authenticated by the authentication server, the SI4091 will verify the *privileges* of the remote user and authorize the appropriate access.

All user privileges, other than those assigned to the Administrator, have to be defined in the RADIUS dictionary. RADIUS attribute 6 which is built into all RADIUS servers defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The following RADIUS attributes are defined for Enterprise NOS user privileges levels:

Table 8. *Enterprise NOS-proprietary Attributes for RADIUS*

User Level/Name	User-Service-Type	Value
User / `user`	*Vendor-supplied*	255
Operator / `oper`	*Vendor-supplied*	252
Administrator / `admin`	*Vendor-supplied*	6
Administrator / `USERID`	*Vendor-supplied*	6

RADIUS Local Access

The administrator has multiple options to allow *local* access via Telnet and SSH. When enabled, local access mode allows the administrator to log in using the `noradius` user name and the regular SI4091 administrator password. Two options are available:

- Secure local access permits remote local access only when the RADIUS server cannot be reached. This option is disabled by default.
- Regular local access permits remote local access at any time, regardless of the RADIUS server status. This option also is disabled by default.

Note: RADIUS local access is always available from the local Console port, regardless of remote local access settings.

Remote local access options can be enabled or disabled using the following CLI commands:

```
SI 4091(config)# [no] radius-server secure-backdoor(Backdoor only when server is down)
SI 4091(config)# [no] radius-server backdoor(Backdoor always available)
```

TACACS+ Authentication

Enterprise NOS supports authentication, authorization, and accounting with networks using the Cisco Systems TACACS+ protocol. The SI4091 functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the SI4091 either through a data or management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP-based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization and accounting.

How TACACS+ Authentication Works

TACACS+ works much in the same way as RADIUS authentication as described on page 90.

1. Remote administrator connects to the SI4091 and provides user name and password.
2. Using Authentication/Authorization protocol, the SI4091 sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using TACACS+ protocol, the authentication server instructs the SI4091 to grant or deny administrative access.

   During a session, if additional authorization checking is needed, the SI4091 checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

TACACS+ Authentication Features in Enterprise NOS

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. Enterprise NOS supports ASCII inbound login to the device. PAP, CHAP and ARAP login methods, TACACS+ change password requests, and one-time password authentication are not supported.

Authorization

Authorization is the action of determining a user's privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization levels and Enterprise NOS management access levels is shown in Table 9. The authorization levels listed in this table must be defined on the TACACS+ server.

Table 9. *Default TACACS+ Authorization Levels*

Enterprise NOS Account Name	TACACS+ Level
`user`	0
`oper`	3
`admin`	6
`USERID` (administrator)	6

Alternate mapping between TACACS+ authorization levels and Enterprise NOS management access levels is shown in Table 10. Use the following command to use the alternate TACACS+ authorization levels:

```
SI 4091(config)# tacacs-server privilege-mapping
```

Table 10. *Alternate TACACS+ Authorization Levels*

Enterprise NOS Account Name	TACACS+ Level
`user`	0–1
`oper`	6–8
`admin`	14–15
`USERID` (administrator)	14–15

If the remote user is successfully authenticated by the authentication server, the SI4091 verifies the *privileges* of the remote user and authorizes the appropriate access.

Local Access

The administrator multiple options to allow *local* access via Telnet and SSH. When enabled, local access allows the administrator to log in using the `notacacs` user name and the regular SI4091 administrator password. Two options are available:

- Secure local access permits remote local access only when the TACACS+ server cannot be reached. This option is disabled by default.
- Regular local access permits remote local access at any time, regardless of the TACACS+ server status. This option also is disabled by default.

Note: TACACS+ local access is always available from the local Console port, regardless of remote local access settings.

Remote local access options can be enabled or disabled using the following CLI commands:

```
SI 4091(config)# [no] tacacs-server secure-backdoor(Backdoor only when server is down)
SI 4091(config)# [no] tacacs-server backdoor     (Backdoor always available)
```

Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization is not performed via TACACS+, there are no TACACS+ accounting messages sent out.

You can use TACACS+ to record and track software login access, configuration changes, and interactive commands.

The SI4091 supports the following TACACS+ accounting attributes:

- `protocol` (`console`/`telnet`/`ssh`)
- `start_time`
- `stop_time`
- `elapsed_time`
- `disc-cause`

Command Authorization and Logging

When TACACS+ Command Authorization is enabled, ENOS configuration commands are sent to the TACACS+ server for authorization. Use the following command to enable TACACS+ Command Authorization:

```
SI 4091(config)# tacacs-server command-authorization
```

When TACACS+ Command Logging is enabled, ENOS configuration commands are logged on the TACACS+ server. Use the following command to enable TACACS+ Command Logging:

```
SI 4091(config)# tacacs-server command-logging
```

The following examples illustrate the format of Enterprise NOS commands sent to the TACACS+ server:

```
authorization request, cmd=shell, cmd-arg=interface ip
accounting request, cmd=shell, cmd-arg=interface ip
authorization request, cmd=shell, cmd-arg=enable
accounting request, cmd=shell, cmd-arg=enable
```

TACACS+ Password Change

Enterprise NOS supports TACACS+ password change. When enabled, users can change their passwords after successful TACACS+ authorization. Use the following command to enable or disable this feature:

```
SI 4091(config)# [no] tacacs-server password-change
```

Use the following commands to change the password for the primary and secondary TACACS+ servers:

```
SI 4091(config)# tacacs-server chpassp     (Change primary TACACS+ password)
SI 4091(config)# tacacs-server chpasss     (Change secondary TACACS+ password)
```

Configuring TACACS+ Authentication on the SI4091

1. Configure the IPv4 addresses of the Primary and Secondary TACACS+ servers, and enable TACACS authentication.

```
SI 4091(config)# tacacs-server primary-host 10.10.1.1
SI 4091(config)# tacacs-server secondary-host 10.10.1.2
SI 4091(config)# tacacs-server enable
```

2. Configure the TACACS+ secret and second secret.

```
SI 4091(config)# tacacs-server primary-host 10.10.1.1 key
<1-32 character secret>
SI 4091(config)# tacacs-server secondary-host 10.10.1.2 key
     <1-32 character secret>
```

CAUTION:
If you configure the TACACS+ secret using any method other than a direct Console port connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default TCP port number used to listen to TACACS+. The well-known port for TACACS+ is 49.

```
SI 4091(config)# tacacs-server port <TCP port number>
```

4. Configure the number of retry attempts, and the timeout period.

```
SI 4091(config)# tacacs-server retransmit 3
SI 4091(config)# tacacs-server timeout 5
```

LDAP Authentication and Authorization

Enterprise NOS supports the LDAP (Lightweight Directory Access Protocol) method to authenticate and authorize remote administrators to manage the SI4091. LDAP is based on a client/server model. The SI4091 acts as a client to the LDAP server. A remote user (the remote administrator) interacts only with the SI4091, not the back-end server and database.

LDAP authentication consists of the following components:

- A protocol with a frame format that utilizes TCP over IP
- A centralized server that stores all the user authorization information
- A client (in this case, the SI4091 is the client)

Each entry in the LDAP server is referenced by its Distinguished Name (DN). The DN consists of the user-account name concatenated with the LDAP domain name. If the user-account name is John, the following is an example DN:

```
uid=John,ou=people,dc=domain,dc=com
```

If the remote user is successfully authenticated by the authentication server, the SI4091 verifies the *privileges* of the remote user and authorizes the appropriate access.

LDAP Backdoor

The administrator has can allow *backdoor* access via Telnet and SSH. When enabled, the backdoor allows the administrator to log in using the `noldap` user name and the regular SI4091 administrator password at any time. This option is disabled by default.

Note: The LDAP backdoor is always available from the local Console port, regardless of the remote backdoor setting.

Remote backdoor access can be enabled or disabled using the following CLI commands:

```
SI 4091(config)# [no] ldap-server backdoor(Backdoor always available)
```

Configuring the LDAP Server

SI4091 user groups and user accounts must reside within the same domain. On the LDAP server, configure the domain to include SI4091 user groups and user accounts, as follows:

- User Accounts:

  Use the *uid* attribute to define each individual user account.

- User Groups:

  Use the *members* attribute in the *groupOfNames* object class to create the user groups. The first word of the common name for each user group must be equal to the user group names defined in the SI4091, as follows:

  - `admin`
  - `oper`
  - `user`

Configuring LDAP Authentication on the SI4091

1. Turn LDAP authentication on, then configure the Primary and Secondary LDAP servers.

```
SI 4091(config)# ldap-server enable
SI 4091(config)# ldap-server primary-host 10.10.1.1
SI 4091(config)# ldap-server secondary-host 10.10.1.2
```

2. Configure the domain name.

```
SI 4091(config)# ldap-server domain <ou=people,dc=my-domain,dc=com>
```

3. If desired, you may change the default TCP port number used to listen to LDAP.

   The well-known port for LDAP is 389.

```
SI 4091(config)# ldap-server port <1-65000>
```

4. Configure the number of retry attempts for contacting the LDAP server, and the timeout period.

```
SI 4091(config)# ldap-server retransmit 3
SI 4091(config)# ldap-server timeout 10
```

5. If desired, change the default LDAP attribute (uid) or add a custom attribute using the following command:

```
SI 4091(config)# ldap-server attribute username <1-128 alphanumeric characters>
```

Chapter 7. Access Control Lists

Access Control Lists (ACLs) are filters that permit or deny traffic for security purposes. They can also be used with QoS to classify and segment traffic in order to provide different levels of service to different traffic types. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

Enterprise NOS 8.4 supports the following ACLs:

- IPv4 ACLs

  Up to 640 ACLs are supported for networks that use IPv4 addressing. IPv4 ACLs are configured using the following CLI command:

```
SI 4091(config)# access-control list <IPv4 ACL number> ?
```

- IPv6 ACLs

  Up to 128 ACLs are supported for networks that use IPv6 addressing. IPv6 ACLs are configured using the following CLI command:

```
SI 4091(config)# access-control list6 <IPv6 ACL number> ?
```

Summary of Packet Classifiers

ACLs allow you to classify packets according to a variety of content in the packet header (such as the source address, destination address, source port number, destination port number, and others). Once classified, packet flows can be identified for more processing.

Regular ACLs allow you to classify packets based on the following packet attributes:

- Ethernet header options (for regular ACLs only)
  - Source MAC address
  - Destination MAC address
  - VLAN number and mask
  - Ethernet type (ARP, IPv4, MPLS, RARP, etc.)
  - Ethernet Priority (the IEEE 802.1p Priority)
- IPv4 header options (for regular ACLs only)
  - Source IPv4 address and subnet mask
  - Destination IPv4 address and subnet mask
  - Type of Service value
  - IP protocol number or name as shown in Table 11:

Table 11. *Well-Known Protocol Types*

Number	Protocol Name
1	icmp
2	igmp
6	tcp
17	udp
89	ospf
112	vrrp

- TCP/UDP header options (for all ACLs)
  - TCP/UDP application source port as shown in Table 12.

Table 12. *Well-Known Application Ports*

Port	TCP/UDP Application	Port	TCP/UDP Application	Port	TCP/UDP Application
20	ftp-data	79	finger	179	bgp
21	ftp	80	http	194	irc
22	ssh	109	pop2	220	imap3
23	telnet	110	pop3	389	ldap
25	smtp	111	sunrpc	443	https
37	time	119	nntp	520	rip
42	name	123	ntp	554	rtsp
43	whois	143	imap	1645/1812	Radius
53	domain	144	news	1813	Radius
69	tftp	161	snmp	1985	Accounting
70	gopher	162	snmptrap		hsrp

  - TCP/UDP application destination port and mask as shown in Table 12.
  - TCP/UDP flag value as shown in Table 13.

Table 13. *Well-Known TCP flag values*

Flag	Value
URG	0x0020
ACK	0x0010
PSH	0x0008
RST	0x0004
SYN	0x0002
FIN	0x0001

- Packet format (for regular ACLs only)
  - Ethernet format (eth2, SNAP, LLC)
  - Ethernet tagging format
  - IP format (IPv4)
- Egress port packets (for all ACLs)

Summary of ACL Actions

Once classified using ACLs, the identified packet flows can be processed differently. For each ACL, an *action* can be assigned. The action determines how the SI4091 treats packets that match the classifiers assigned to the ACL. SI4091 ACL actions include the following:

- Pass or Drop the packet
- Re-mark the packet with a new DiffServ Code Point (DSCP)
- Re-mark the 802.1p field
- Set the COS queue

Assigning Individual ACLs to a Port

Once you configure an ACL, you must assign the ACL to the appropriate ports. Each port can accept multiple ACLs, and each ACL can be applied for multiple ports. ACLs can be assigned individually, or in groups.

To assign an individual ACLs to a port, use the following IP Interface Mode commands:

```
SI 4091(config)# interface port <port>
SI 4091(config-ip)# access-control list <IPv4 ACL number>
SI 4091(config-ip)# access-control list6 <IPv6 ACL number>
```

When multiple ACLs are assigned to a port, higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs. ACL order of precedence is discussed in the next section.

To create and assign ACLs in groups, see "ACL Groups" on page 107.

ACL Order of Precedence

When multiple ACLs are assigned to a port, they are evaluated in numeric sequence, based on the ACL number. Lower-numbered ACLs take precedence over higher-numbered ACLs. For example, ACL 1 (if assigned to the port) is evaluated first and has top priority.

If multiple ACLs match the port traffic, only the action of the one with the lowest ACL number is applied. The others are ignored.

The ACL number is the sole factor in determining ACL order of precedence. The order in which ACLs are applied to a port does not affect the order of precedence, nor does the ACL Group number (see "ACL Groups" on page 107), the order in which an ACL is assigned to an ACL Group, or the order in which the ACL Group is assigned to a port.

If no assigned ACL matches the port traffic, no ACL action is applied.

ACL Groups

To assist in organizing multiple ACLs and assigning them to ports, you can place ACLs into ACL Groups, thereby defining complex traffic profiles. ACLs and ACL Groups can then be assigned on a per-port basis. Any specific ACL can be assigned to multiple ACL Groups, and any ACL or ACL Group can be assigned to multiple ports. If, as part of multiple ACL Groups, a specific ACL is assigned to a port multiple times, only one instance is used. The redundant entries are ignored.

- **Individual ACLs**
  The SI4091 supports up to 640 ACLs. Each ACL defines one filter rule for matching traffic criteria. Each filter rule can also include an action (permit or deny the packet). For example:

ACL 1: VLAN = 1 SIP = 10.10.10.1 (255.255.255.0) Action = permit

- **Access Control List Groups**
  An Access Control List Group (ACL Group) is a collection of ACLs. For example:

ACL Group 1
ACL 1: VLAN = 1 SIP = 10.10.10.1 (255.255.255.0) Action = permit
ACL 2: VLAN = 2 SIP = 10.10.10.2 (255.255.255.0) Action = deny
ACL 3: Priority = 7 DIP = 10.10.10.3 (255.255.255.0) Action = permit

  ACL Groups organize ACLs into traffic profiles that can be more easily assigned to ports. The SI4091 supports up to 640 ACL Groups.

Note: ACL Groups are used for convenience in assigning multiple ACLs to ports. ACL Groups have no effect on the order in which ACLs are applied (see "ACL Order of Precedence" on page 106). All ACLs assigned to the port (whether individually assigned or part of an ACL Group) are considered as individual ACLs for the purposes of determining their order of precedence.

Assigning ACL Groups to a Port

To assign an ACL Group to a port, use the following commands:

```
SI 4091(config)# interface port <port>
SI 4091(config-ip)# access-control group <ACL group number>
```

ACL Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the SI4091 by configuring a QoS meter (if desired) and assigning ACLs to ports.

Note: When you add ACLs to a port, make sure they are ordered correctly in terms of precedence (see "ACL Order of Precedence" on page 106).

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- **In-Profile**–If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile**–If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kbps (1000 bits per second in each Kbps). All traffic within this Committed Rate is In-Profile. Additionally, you can set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level that traffic should receive.
- Change the 802.1p priority of a packet.

ACL Port Mirroring

For regular ACLs, packets that match an ACL on a specific port can be mirrored to another SI4091 port for network diagnosis and monitoring.

The source port for the mirrored packets cannot be a portchannel, but may be a member of a portchannel.

The destination port to which packets are mirrored must be a physical port.

ACL mirroring cannot be used on ACLs on which an action (such as permit or deny) is configured.

Use the following command to add mirroring to an ACL:

```
SI 4091(config)# access-control list <ACL number> mirror port <destination port>
```

The ACL must be also assigned to it target ports as usual (see "Assigning Individual ACLs to a Port" on page 105, or "Assigning ACL Groups to a Port" on page 108).

Viewing ACL Statistics

ACL statistics display how many packets have “hit” (matched) each ACL. Use ACL statistics to check filter performance or to debug the ACL filter configuration.

You must enable statistics for each ACL that you wish to monitor:

```
SI 4091(config)# access-control list <ACL number> statistics
```

ACL Logging

ACLs are generally used to enhance port security. Traffic that matches the characteristics (source addresses, destination addresses, packet type, etc.) specified by the ACLs on specific ports is subject to the actions (chiefly permit or deny) defined by those ACLs. Although switch statistics show the number of times particular ACLs are matched, the ACL logging feature can provide additional insight into actual traffic patterns on the switch, providing packet details in the system log for network debugging or security purposes.

Enabling ACL Logging

By default, ACL logging is disabled. Enable or disable ACL logging on a per-ACL basis as follows:

```
SI 4091(config)# [no] access-control list <IPv4 ACL number> log
SI 4091(config)# [no] access-control list6 <IPv6 ACL number> log
```

Logged Information

When ACL logging is enabled on any particular ACL, the switch will collect information about packets that match the ACL. The information collected depends on the ACL type:

- For IP-based ACLs, information is collected regarding
  - Source IP address
  - Destination IP address
  - TCP/UDP port number
  - ACL action
  - Number of packets logged

  For example:

  ```
  Sep 27  4:20:28 DUT3 NOTICE  ACL-LOG: %IP ACCESS LOG: list
  ACL-IP-12-IN denied tcp 1.1.1.1 (0) -> 200.0.1.2 (0), 150
  packets.
  ```

- For MAC-based ACLs, information is collected regarding
  - Source MAC address
  - Source IP address
  - Destination IP address
  - TCP/UDP port number
  - ACL action
  - Number of packets logged

  For example:

  ```
  Sep 27  4:25:38 DUT3 NOTICE  ACL-LOG: %MAC ACCESS LOG: list
  ACL-MAC-12-IN permitted tcp 1.1.1.2 (0) (12,
  00:ff:d7:66:74:62) -> 200.0.1.2 (0) (00:18:73:ee:a7:c6), 32
  packets.
  ```

Rate Limiting Behavior

Because ACL logging can be CPU-intensive, logging is rate-limited. By default, the switch will log only 10 matching packets per second. This pool is shared by all log-enabled ACLs. The global rate limit can be changed as follows:

```
SI 4091(config)# access-control log rate-limit <1-1000>
```

Where the limit is specified in packets per second.

Log Interval

For each log-enabled ACL, the first packet that matches the ACL initiates an immediate message in the system log. Beyond that, additional matches are subject to the log interval. By default, the switch will buffer ACL log messages for a period of 300 seconds. At the end of that interval, all messages in the buffer are written to the system log. The global interval value can be changed as follows:

```
SI 4091(config)# access-control log interval <5-600>
```

Where the interval rate is specified in seconds.

In any given interval, packets that have identical log information are condensed into a single message. However, the packet count shown in the ACL log message represents only the logged messages, which due to rate-limiting, may be significantly less than the number of packets actually matched by the ACL.

Also, the switch is limited to 64 different ACL log messages in any interval. Once the threshold is reached, the oldest message will be discarded in favor of the new message, and an overflow message will be added to the system log.

ACL Logging Limitations

ACL logging reserves packet queue 1 for internal use. Features that allow remapping packet queues (such as CoPP) may not behave as expected if other packet flows are reconfigured to use queue 1.

ACL Configuration Examples

The following are sample ACL configurations.

ACL Example 1

Use this configuration to block traffic to a specific host. All traffic that ingresses on port EXT1 is denied if it is destined for the host at IP address 100.10.1.1

1. Configure an Access Control List.

```
SI 4091(config)# access-control list 1 ipv4 destination-ip-address
    100.10.1.1
SI 4091(config)# access-control list 1 action deny
```

2. Add ACL 1 to port EXT1.

```
SI 4091(config)# interface port ext1
SI 4091(config-if)# access-control list 1
SI 4091(config-if)# exit
```

ACL Example 2

Use this configuration to block traffic from a network destined for a specific host address. All traffic that ingresses in port EXT2 with source IP from class 100.10.1.0/24 and destination IP 200.20.2.2 is denied.

1. Configure an Access Control List.

```
SI 4091(config)# access-control list 2 ipv4 source-ip-address 100.10.1.0
    255.255.255.0
SI 4091(config)# access-control list 2 ipv4 destination-ip-address
    200.20.2.2 255.255.255.255
SI 4091(config)# access-control list 1 action deny
```

2. Add ACL 2 to port EXT2.

```
SI 4091(config)# interface port ext2
SI 4091(config-if)# access-control list 2
SI 4091(config-if)# exit
```

ACL Example 3

This configuration blocks traffic from a network that is destined for a specific egress port. All traffic that ingresses port EXT1 from the network 100.10.1.0/24 and is destined for port 3 is denied.

1. Configure an Access Control List.

```
SI 4091(config)# access-control list 4 ipv4 source-ip-address 100.10.1.0
    255.255.255.0
SI 4091(config)# access-control list 4 egress-port 3
SI 4091(config)# access-control list 4 action deny
```

2. Add ACL 4 to port EXT1.

```
SI 4091(config)# interface port ext1
SI 4091(config-if)# access-control list 4
SI 4091(config-if)# exit
```

Using Storm Control Filters

Excessive transmission of broadcast or multicast traffic, such as during denial-of-service (DoS) attacks, can result in a broadcast storm, degrading network performance. Common symptoms of a broadcast storm are slow network response times and network operations timing out.

Unicast packets whose destination MAC address is not in the Forwarding Database are *unknown unicasts.* When an unknown unicast is encountered, the switch normally handles it like a broadcast packet and floods it to all other ports in the VLAN (broadcast domain). A high rate of unknown unicast traffic can have the same negative effects as a broadcast storm.

To defend against the effects of network storms, the SI4091 provides filters that can limit the rate of packet transmissions for the specific types of traffic associated with these types of storms:st packets

- Multicast packets
- Unknown unicast packets (destination lookup failure)

Configure broadcast filters on each port that requires broadcast storm control. Set a threshold that defines the total number of broadcast packets transmitted (0–2097151) per second. When the threshold is reached, no more packets of the specified type are transmitted by the port.

To filter broadcast packets on a port, use the following commands:

```
SI 4091(config)# interface port 1
SI 4091(config-if)# stormcontrol broadcast level pps <packets per second>
```

To filter multicast packets on the port, use the following commands:

```
SI 4091(config-if)# stormcontrol multicast level pps <packets per second>
```

To filter unknown unicast packets on the port, use the following commands:

```
SI 4091(config-if)# stormcontrol unicast level pps <packets per second>
SI 4091(config-if)# exit
```

Part 3: Basic Features

Chapter 8. External Uplink Interfaces

Internal server ports and external uplink interfaces can each be associated with one or more VLANs. Traffic in the switch is segregated into distinct domains based on these VLAN memberships; Traffic on the ports of one VLAN is isolated from traffic on the others. This permits multiple domains to co-exist securely in the switch.

The switch automatically prevents data loops among connected links without relying on slow, complicated IEEE 802.1d Spanning-Tree Protocols. Within the switch, loop-free operation is enforced by permitting only one external uplink interface for each VLAN.

Figure 3. External Uplink Options

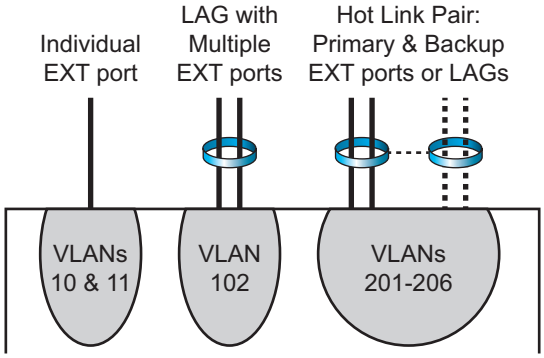

Each external interface can be comprised of one of the following:

- A single physical external port.
- Multiple external ports that are aggregated together into a single Link Aggregation Group (LAG). For more, see Chapter 10, "Ports and Link Aggregation."
- A primary and backup pair of Hot Links entities, each of which can in turn be independently comprised of a single port or a LAG. For more, see "Hot Links" on page 237.

Each VLAN must have one and only one external uplink interface. However, any specific uplink interface may be used by multiple VLANs.

By default, all internal and external non-management ports are assigned to VLAN 1, and all external ports are aggregated together into one LAG (with LACP key 1000, Portchannel ID 25) as a single uplink interface.

Uplink Interface Configuration

The general configuration process for the SI4091 is as follows:

1. Create a VLAN for each desired domain.
2. Create one uplink entity for each domain that will use uplink LAGs or Hot Links.
3. Assign internal ports and uplinks to each VLAN

The following examples provide specifics.

Example 1: Transparent VLANs

In this example, the switch is configured for transparent (VLAN-agnostic) behavior.

In this configuration, the SI4091 connects the internal (server) ports INT1–INT14 to multiple upstream domains on the external LAG of EXT1–EXT10, and no individual VLAN configuration is required on the SI4091. All the internal ports have access to all VLANs, and broadcasts for any VLAN are forwarded to all internal ports. All individual VLAN participation, port filters, etc. are expected to be configured on the connected servers and upstream devices.

1. Create a VLAN for each desired domain.

```
SI 4091(config)# vlan 4081
SI 4091(config-vlan)# name “Transparent”
SI 4091(config-vlan)# exit
```

2. Create one uplink entity for each domain that will use uplink LAGs.

   In this case, by default, EXT1 though EXT10 are configured in a dynamic LAG (LACP key 1000, Portchannel ID 25), so no additional LAG configuration is required. However, for example purposes, the LAG could be manually configured as follows:

```
SI 4091(config)# portchannel 25 lacp key 1000
SI 4091(config)# interface port ext1-ext10
SI 4091(config-if)# lacp key 1000
SI 4091(config-if)# lacp mode active
SI 4091(config-if)# exit
```

3. Assign the desired internal ports and uplink to the VLAN.

```
SI 4091(config)# interface port int1-int14(Internal ports)
SI 4091(config-if)# switchport mode access
SI 4091(config-if)# switchport access vlan 4081
SI 4091(config-if)# tagpvid-ingress
SI 4091(config-if)# exit
SI 4091(config)# interface portchannel lacp 1000(External uplink)
SI 4091(config-if)# switchport mode access
SI 4091(config-if)# switchport access vlan 4081
SI 4091(config-if)# tagpvid-ingress
SI 4091(config-if)# exit
```

4. You can confirm the resulting configuration using the `show running-config` command.

Example 2: Multiple VLAN Domains

In this example, instead of one VLAN-transparent domain, the switch enforces separation among multiple upstream VLAN domains.

Figure 4. Multiple VLAN Domains

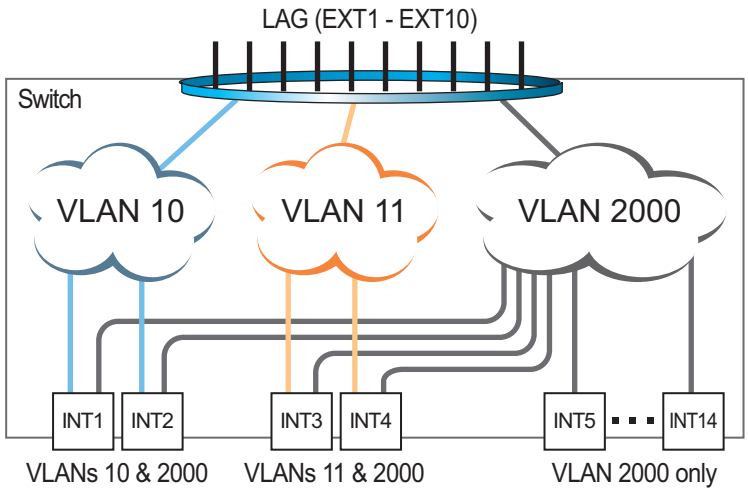

As shown in Figure 4, all internal (server) ports have access to main VLAN 2000. Two of the internal ports (INT1 and INT2) also have access to VLAN 10, and two others (INT3 and INT4) have access to VLAN 11.

In this scenario, individual VLAN configuration is required on the SI4091, as well as on connected servers and upstream devices.

1. Create VLANs for each desired domain:

```
SI 4091(config)# vlan 2000                (Main VLAN for all ports)
SI 4091(config-vlan)# exit

SI 4091(config)# vlan 10                  (VLAN for INT1 and INT2)
SI 4091(config-vlan)# exit

SI 4091(config)# vlan 11                  (VLAN for INT3 and INT4)
SI 4091(config-vlan)# exit
```

2. Create the desired uplink entity for each required LAG.

   In this case, by default, EXT1 though EXT10 are configured in a dynamic LAG (LACP key 1000, Portchannel ID 25), so no additional LAG configuration is required. However, for example purposes, the LAG could be manually configured as follows:

```
SI 4091(config)# portchannel 25 lacp key 1000
SI 4091(config)# interface port ext1-ext10
SI 4091(config-if)# lacp key 1000
SI 4091(config-if)# lacp mode active
SI 4091(config-if)# exit
```

3. Assign the desired internal ports and uplink to their VLANs.

```
SI 4091(config)# interface port int1-int2
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 10,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port int3-int4
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 11,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port int5-int14
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port ext1-ext10
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 10,11,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit
```

4. You can confirm the resulting configuration using the `show running-config` command.

Example 3: Multiple VLAN Domains and FCoE

If the prior example on page 122 were to also require FCoE on VLAN 1002 (the typical FCoE VLAN), the procedure would be modified as follows:

1. Create VLANs for each desired domain:

```
SI 4091(config)# vlan 2000                (Main VLAN for all ports)
SI 4091(config-vlan)# exit

SI 4091(config)# vlan 10                  (VLAN for INT1 and INT2)
SI 4091(config-vlan)# exit

SI 4091(config)# vlan 11                  (VLAN for INT3 and INT4)
SI 4091(config-vlan)# exit

SI 4091(config)# vlan 1002                (VLAN for FCoE SAN)
SI 4091(config-vlan)# exit
```

2. Create the desired uplink entity for each required LAG if required:

```
SI 4091(config)# portchannel 25 lacp key 1000
SI 4091(config)# interface port ext1-ext10
SI 4091(config-if)# lacp key 1000
SI 4091(config-if)# lacp mode active
SI 4091(config-if)# exit
```

3. Assign the desired internal ports and uplink to their VLANs.

```
SI 4091(config)# interface port int1-int2
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 10,1002,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port int3-int4
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 11,1002,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port int5-int14
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 1002,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit

SI 4091(config)# interface port ext1-ext10
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan 10,11,1002,2000
SI 4091(config-if)# switchport trunk native vlan 2000
SI 4091(config-if)# exit
```

4. Ensure that CEE and FIP Snooping are enabled:

```
SI 4091(config)# cee enable
SI 4091(config)# fcoe fips enable
```

5. You can confirm the resulting configuration using the `show running-config` command.

Chapter 9. VLANs

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs).

VLANs are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments. The following topics are discussed in this chapter:

VLANs Overview

Setting up VLANs is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each SI4091 port connects to a segment that is a single broadcast domain. When a port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

The SI4091 automatically supports jumbo frames with a Maximum Transmission Unit (MTU) of 9,216 bytes. Within each frame, 18 bytes are reserved for the Ethernet header and CRC trailer. The remaining space in the frame (up to 9,198 bytes) comprise the packet, which includes the payload of up to 9,000 bytes and any additional overhead, such as 802.1q or VLAN tags. Jumbo frame support is enabled by default, requires no manual configuration, and cannot be manually disabled.

Note: Jumbo frames are not supported for traffic sent to SI4091 management interfaces.

VLANs and Port VLAN ID Numbers

VLAN Numbers

The SI4091 supports up to 4095 VLANs. The following VLANs are reserved:

- VLAN 1 is the default VLAN for the internal and external ports.
- VLAN 1002 is the default VLAN for Fibre Channel over Ethernet (FCoE) traffic. The FCoE VLAN may be changed, or if not used for FCoE traffic, it may be assigned for other purposes.
- VLAN 4091 is known as a *black-hole* VLAN since it permits no egress traffic. By default, VLAN 4091 is empty.
- VLAN 4095 is reserved for use by the management network, which includes the management ports (MGMT1 and EXTM).

  Management functions can also be assigned to other VLANs (using the following command:

```
SI 4091(config)# vlan <x>
SI 4091(config-vlan)# management
```

Use the following command to view VLAN information:

```
SI 4091(config)# show vlan
VLAN                Name                Status MGT           Ports
----  --------------------------------  ------ ---
------------------------
1     Default VLAN                      ena    dis  INT1-EXT10
4091  Black-hole VLAN                   ena    dis  empty
4095  Mgmt VLAN                         ena    ena  EXTM MGT1

Primary  Secondary  Type           Ports
-------  ---------  ---------------
----------------------------------------
```

Note: The sample screens that appear in this document might differ slightly from the screens displayed by your system. Screen content varies based on the type of chassis unit that you are using and the firmware versions and options that are installed.

PVID/Native VLAN Numbers

Each port in the SI4091 has a configurable default VLAN number, known as its *PVID* or native VLAN. This is the assumed VLAN number for all port traffic not explicitly tagged with its own VLAN number.

By default, the PVID correlates to the default VLAN ID (see "VLAN Numbers" on page 129). The PVID for each port can be configured to any VLAN number between 1 and 4094.

Use the following CLI commands to view PVIDs:

- To display port information—

```
SI 4091# show interface information
   (or)
SI 4091# show interface trunk
Alias  Port Tag RMON Lrn Fld PVID    DESCRIPTION             VLAN(s)
            Trk              NVLAN
------ ---- --- ---- --- --- ------ ------------
----------------------------
INT1   1     n   d    e   e     1#  INT1           1
INT2   2     n   d    e   e     1#  INT2           1
INT3   3     n   d    e   e     1#  INT3           1
INT4   4     n   d    e   e     1#  INT4           1
INT5   5     n   d    e   e     1#  INT5           1
INT6   6     n   d    e   e     1#  INT6           1
INT7   7     n   d    e   e     1#  INT7           1
INT8   8     n   d    e   e     1#  INT8           1
INT9   9     n   d    e   e     1#  INT9           1
INT10  10    n   d    e   e     1#  INT10          1
INT11  11    n   d    e   e     1#  INT11          1
INT12  12    n   d    e   e     1#  INT12          1
INT13  13    n   d    e   e     1#  INT13          1
INT14  14    n   d    e   e     1#  INT14          1
EXT1   15    n   d    e   e     1#  EXT1           1
EXT2   16    n   d    e   e     1#  EXT2           1
EXT3   17    n   d    e   e     1#  EXT3           1
EXT4   18    n   d    e   e     1#  EXT4           1
EXT5   19    n   d    e   e     1#  EXT5           1
EXT6   20    n   d    e   e     1#  EXT6           1
EXT7   21    n   d    e   e     1#  EXT7           1
EXT8   22    n   d    e   e     1#  EXT8           1
EXT9   23    n   d    e   e     1#  EXT9           1
EXT10  24    n   d    e   e     1#  EXT10          1
EXTM   25    n   d    e   e  4095   EXTM           4095
MGT1   26    y   d    e   e  4095   MGT1           4095

* = PVID/Native-VLAN is tagged.
# = PVID is ingress tagged.
Trk  = Trunk mode
NVLAN = Native-VLAN
```

Note: The sample output that appears in this document might differ slightly from that displayed by your system. Output varies based on the type of blade chassis unit that you are using and the firmware versions and options that are installed.

- To configure access mode ports:

```
SI 4091(config)# interface port <port number>
SI 4091(config-if)# switchport access vlan <VLAN ID>
SI 4091(config-if)# exit
```

- To configure trunk mode ports:

```
SI 4091(config)# interface port <port number>
SI 4091(config-if)# switchport mode trunk
SI 4091(config-if)# switchport trunk allowed vlan <VLAN ID>
SI 4091(config-if)# exit
```

Note: The command `switchport mode trunk` must be entered before the command `switchport trunk allowed vlan` <*VLAN ID*>. If the latter command is entered first, an error will be generated.

Each port on the SI4091 can belong to one or more VLANs, and each VLAN can have any number of ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN *tagging* enabled (see "VLAN Tagging/Trunk Mode" on page 132).

Black-Hole VLAN

To prevent external ports from inadvertently becoming active immediately when they are connected or moved to a new VLAN, assign unused external ports to VLAN 4091. VLAN 4091 is known as a system *black-hole* VLAN since egress traffic is not permitted. Members of the black-hole VLAN do not have access to other ports.

Although VLAN 4091 is the default for use as the system black-hole, any VLAN not reserved for other purposes can be used instead.

To reserve the system black-hole VLAN, use the following Privileged EXE command in the CLI:

```
SI 4091(config)# system internal-vlan <reserved VLAN range>
```

To set the system black-hole VLAN, use the following Privileged EXEC command in the CLI:

```
SI 4091(config)# system black-hole vlan <VLAN ID 2–4094>
```

VLAN Tagging/Trunk Mode

Enterprise NOS software supports 802.1Q VLAN *tagging*, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header of a packet, allowing each port to belong to multiple VLANs. When you add a port to multiple VLANs, you also must enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID = 3) assigns all untagged frames received on this port to VLAN 3. Any untagged frames received by the SI4091 are classified with the PVID of the receiving port.
- Tagged frame—a frame that carries VLAN tagging information in the header. This VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the SI4091 through a port that is configured as a tagged port.
- Untagged frame— a frame that does not carry any VLAN tagging information in the frame header.
- Untagged member—a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the SI4091 through an untagged member port, the frame header remains unchanged. When a tagged frame exits the SI4091 through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the SI4091 through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the SI4091 through a tagged member port, the frame header remains unchanged (original VID remains).

Figure 5. Default VLAN settings

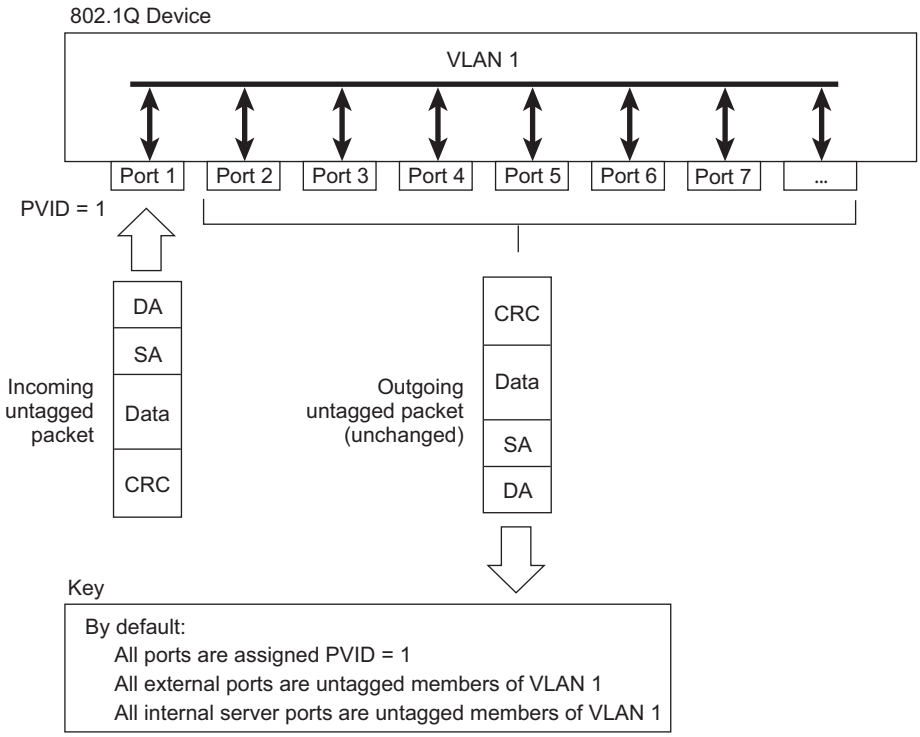

Note: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your SI4091 model.

When a VLAN is configured, ports are added as members of the VLAN, and the ports are defined as either *tagged* or *untagged* (see Figure 6 through Figure 9).

By default, VLAN configuration for all SI4091 internal and external ports is set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in Figure 5, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID = 1).

Figure 6 through Figure 9 illustrate generic examples of VLAN tagging.

In Figure 6, untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Note: The port assignments in the following figures are general examples and are not meant to match any specific SI4091.

Figure 6. Port-based VLAN assignment

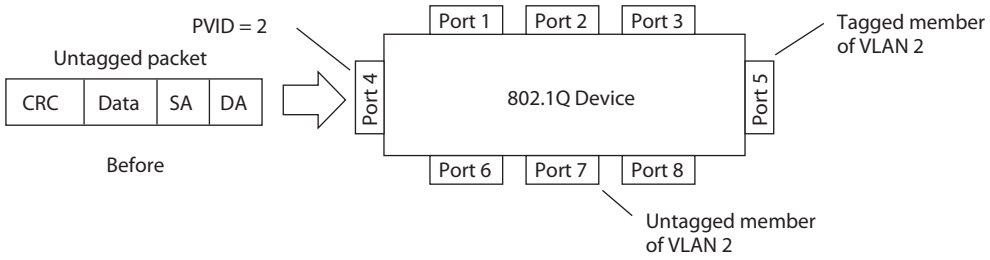

As shown in Figure 7, the untagged packet is marked (tagged) as it leaves the SI4091 through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the SI4091 through port 7, which is configured as an untagged member of VLAN 2.

Figure 7. 802.1Q tagging (after port-based VLAN assignment)

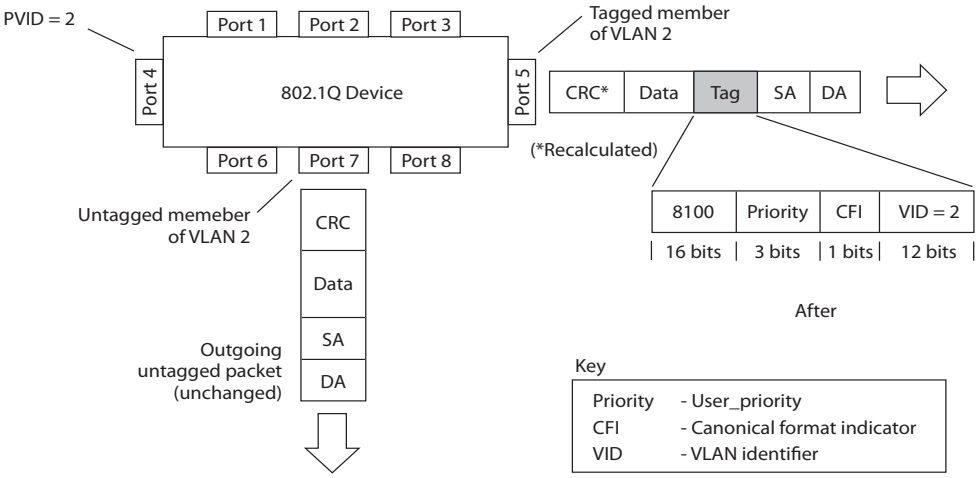

In Figure 8, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 8. 802.1Q tag assignment

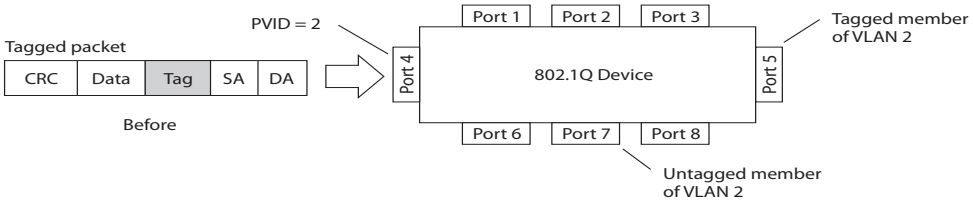

As shown in Figure 9, the tagged packet remains unchanged as it leaves the SI4091 through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the SI4091 through port 7, which is configured as an untagged member of VLAN 2.

Figure 9. 802.1Q tagging (after 802.1Q tag assignment)

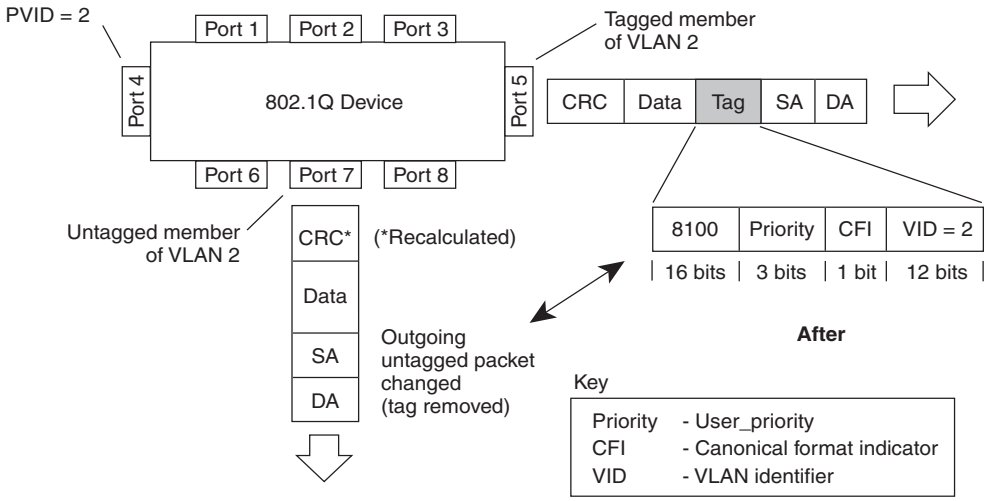

Note: Set the configuration to factory default to return all non-management ports to their original VLANs:

```
SI 4091(config)# boot configuration-block factory
```

Ingress VLAN Tagging

Ingress VLAN tagging is enabled on all ingress ports by default, but can be disabled or enabled on a per-port basis. When a packet is received on an ingress port, and if ingress tagging is enabled on the port, a VLAN tag with the port PVID is inserted into the packet as the outer VLAN tag. Depending on the egress port setting (tagged or untagged), the outer tag of the packet is retained or removed when it leaves the egress port.

Ingress VLAN tagging is used to tunnel packets through a public domain without altering the original 802.1Q status.

When ingress tagging is enabled on a port, all packets, whether untagged or tagged, will be tagged again. As shown in Figure 10, when tagging is enabled on the egress port, the outer tag of the packet is retained when it leaves the egress port. If tagging is disabled on the egress port, the outer tag of the packet is removed when it leaves the egress port.

Figure 10. 802.1Q tagging (after ingress tagging assignment)

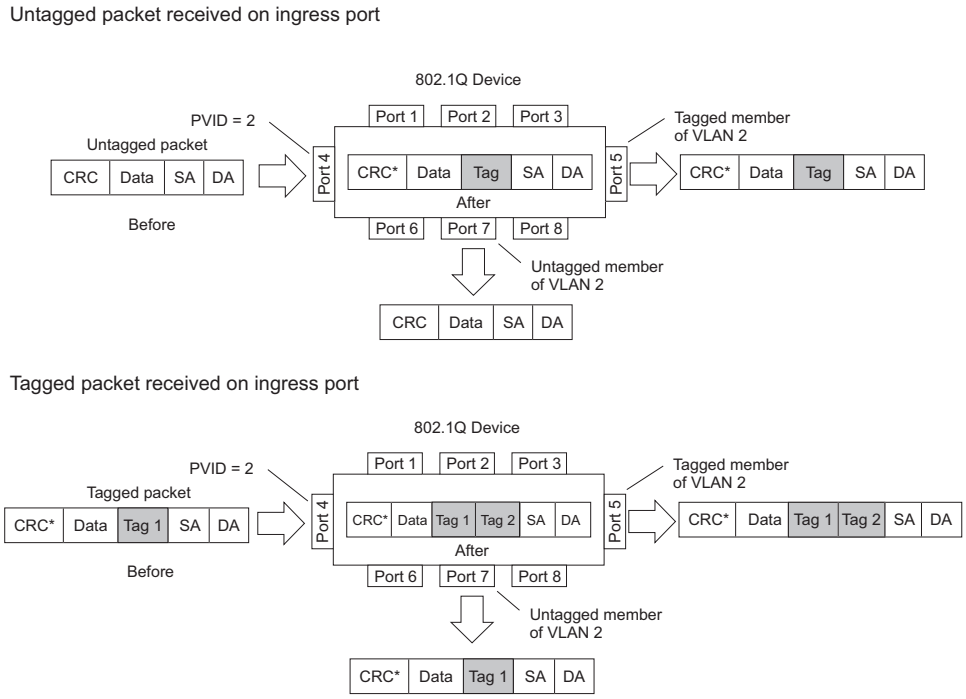

By default, ingress tagging is enabled on all non-management ports. To disable or re-enable ingress tagging on a port, use the following commands:

```
SI 4091(config)# interface port <number>
SI 4091(config-if)# [no] tagpvid-ingress
```

Note: Ingress tagging cannot be enabled on management ports.

VLAN Topologies and Design Considerations

- By default, tagging is disabled for all internal ports and all external ports.
- By default, internal and external ports are members of VLAN 1.
- Multiple management VLANs can be configured on the SI4091, in addition to the default VLAN 4095, using the following commands:

```
SI 4091(config)# vlan <x>
SI 4091(config-vlan)# management
```

Example: Multiple VLANs with Tagging Adapters

Figure 11. Multiple VLANs with VLAN-Tagged Gigabit Adapters

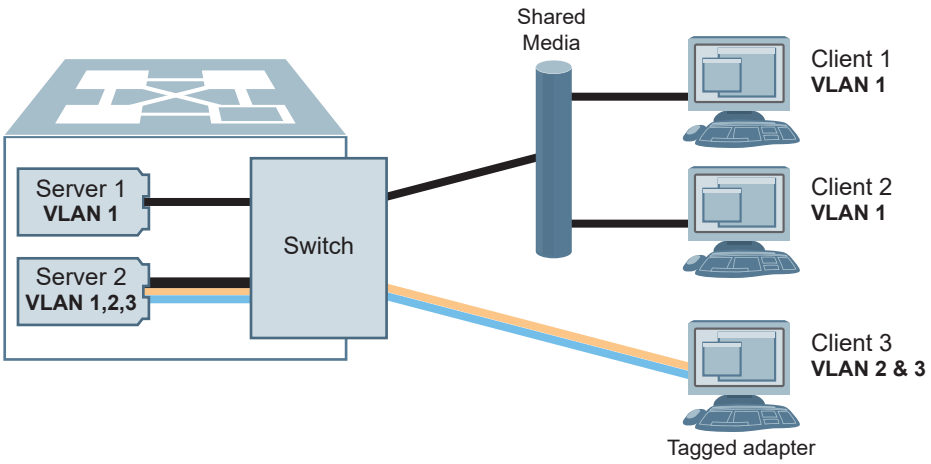

The features of this VLAN are described in the following table:

Component	Description
Switch	This SI4091 is configured for three VLANs that represent three different IP subnets. Two servers and three clients are attached to the SI4091.
Server 1	This server is a member of VLAN 1 and has presence in only one IP subnet. The associated internal SI4091 port is only a member of VLAN 1, so tagging is disabled.
Server 2	This high-use server needs to be accessed from all VLANs and IP subnets. The server has a VLAN-tagging adapter installed with VLAN tagging turned on. The adapter is attached to one of the internal SI4091 ports. The port is a member of VLANs 1, 2, and 3, and has tagging enabled. Because of the VLAN tagging capabilities of both the adapter and the SI4091, the server is able to communicate on all three IP subnets in this network while broadcast separation between all three VLANs and subnets maintained.
Client 1 & 2	These client PCs are attached to a shared media hub that is then connected to the SI4091. They belong only to VLAN 1 and are logically in the same IP subnet as Server 1 & 2. The associated external SI4091 port has tagging disabled.
Client 3	A member of both VLAN 2 and VLAN 3, this client PC has a VLAN-tagging Gigabit Ethernet adapter installed. It can communicate with Server 2. The associated external SI4091 port is a member of VLAN 2 and VLAN 3, and has tagging enabled.

Note: VLAN tagging is required only on ports that are connected to other network devices that connect to tag-capable end-stations, such as servers with VLAN-tagging adapters.

Private VLANs

Private VLANs provide Layer 2 isolation between the ports within the same broadcast domain. Private VLANs can control traffic within a VLAN domain, and provide port-based security for host servers.

Enterprise NOS supports Private VLAN configuration as described in RFC 5517.

Use Private VLANs to partition a VLAN domain into sub-domains. Each sub-domain is comprised of one primary VLAN and one secondary VLAN, as follows:

- Primary VLAN—carries unidirectional traffic downstream from promiscuous ports. Each Private VLAN has only one primary VLAN. All ports in the Private VLAN are members of the primary VLAN.
- Secondary VLAN—Secondary VLANs are internal to a private VLAN domain, and are defined as follows:
  - Isolated VLAN—carries unidirectional traffic upstream from the host servers toward ports in the primary VLAN. Each Private VLAN can contain only one Isolated VLAN.
  - Community VLAN—carries upstream traffic from ports in the community VLAN to other ports in the same community, and to ports in the primary VLAN. Each Private VLAN can contain multiple community VLANs.

After you define the primary VLAN and one or more secondary VLANs, you map the secondary VLAN(s) to the primary VLAN.

Private VLAN Ports

Private VLAN ports are defined as follows:

- Promiscuous—A promiscuous port is an external port that belongs to the primary VLAN. The promiscuous port can communicate with all the interfaces, including ports in the secondary VLANs (Isolated VLAN and Community VLANs).
- Isolated—An isolated port is a host port that belongs to an isolated VLAN. Each isolated port has complete layer 2 separation from other ports within the same private VLAN (including other isolated ports), except for the promiscuous ports.
  - Traffic sent to an isolated port is blocked by the Private VLAN, except the traffic from promiscuous ports.
  - Traffic received from an isolated port is forwarded only to promiscuous ports.
- Community—A community port is a host port that belongs to a community VLAN. Community ports can communicate with other ports in the same community VLAN, and with promiscuous ports. These interfaces are isolated at layer 2 from all other interfaces in other communities and from isolated ports within the Private VLAN.

Configuration Guidelines

The following guidelines apply when configuring Private VLANs:

- Management VLANs cannot be Private VLANs. Management ports cannot be members of a Private VLAN.
- The default VLAN 1 cannot be a Private VLAN.
- IGMP Snooping must be disabled on Private VLANs.
- A VLAN pair is a primary VLAN and one associated secondary VLAN (isolated or community). The maximum number of VLAN pairs per port is 16.

Configuration Example

Follow this procedure to configure a Private VLAN.

1. Select a VLAN and define the Private VLAN type as primary.

```
SI 4091(config)# vlan 700
SI 4091(config-vlan)# private-vlan primary
SI 4091(config-vlan)# exit
```

2. Configure a promiscuous port for VLAN 700.

```
SI 4091(config)# interface port int1
SI 4091(config-if)# switchport mode private-vlan
SI 4091(config-if)# switchport private-vlan mapping 700
SI 4091(config-if)# exit
```

3. Configure two secondary VLANs: isolated VLAN and community VLAN.

```
SI 4091(config)# vlan 701
SI 4091(config-vlan)# private-vlan isolated
SI 4091(config-vlan)# exit
SI 4091(config)# vlan 702
SI 4091(config-vlan)# private-vlan community
SI 4091(config-vlan)# exit
```

4. Map secondary VLANs to primary VLAN.

```
SI 4091(config)# vlan 700
SI 4091(config-vlan)# private-vlan association 701,702
SI 4091(config-vlan)# exit
```

5. Configure host ports for secondary VLANs.

```
SI 4091(config)# interface port int2
SI 4091(config-if)# switchport mode private-vlan
SI 4091(config-if)# switchport private-vlan host-association 700 701
SI 4091(config-if)# exit

SI 4091(config)# interface port int3
SI 4091(config-if)# switchport mode private-vlan
SI 4091(config-if)# switchport private-vlan host-association 700 702
SI 4091(config-if)# exit
```

6. Verify the configuration.

```
SI 4091(config)# show vlan private-vlan

Primary  Secondary      Type               Ports
-------  --------- --------------- --------------------------------
700         701        isolated            INT1-INT2
700         702        community           1NT1 INT3
```

Configuration Considerations

- When you first configure an uplink interface (port or portchannel) in trunk mode and then try to add it to a private vlan, an error is returned. To get around this, first add the uplink interface into the private vlan in trunk mode and *then* configure it in trunk mode.
- If you have a switch with a private VLAN domain with a static portchannel that has one member in each VLAN of the private VLAN domain, if you try to delete the configured static portchannel, an error message will display for each external port.These error messages can be ignored.

Chapter 10. Ports and Link Aggregation

Link Aggregation Groups (LAGs) can provide super-bandwidth, multi-link connections between the SI4091 and other LAG-capable devices. A LAG is a group of ports that act together, combining their bandwidth to create a single, larger virtual link. This chapter provides configuration background and examples for aggregating multiple ports together:

Aggregation Overview

When using port LAGs between two devices, as shown in Figure 12, you can create a virtual link between them, operating with combined throughput levels that depends on how many physical ports are included.

Two types of aggregation are available: static LAGs (portchannels) and dynamic LACP LAGs. Up to 52 LAGs of each type are supported, depending of the number and type of available ports. Each LAG can include up to 24 member ports.

Figure 12. LAG Group

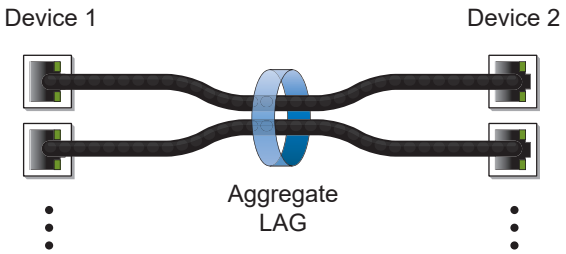

LAGs are also useful for connecting a SI4091 to Lenovo switches and third-party devices that support link aggregation, such as Cisco routers and switches with EtherChannel technology (*not* ISL aggregation technology) and Sun's Quad Fast Ethernet Adapter. LAG technology is compatible with these devices when they are configured manually.

LAG traffic is statistically distributed among the ports in a LAG, based on a variety of configurable options.

Also, since each LAG is comprised of multiple physical links, the LAG is inherently fault tolerant. As long as one connection between the devices is available, the LAG remains active and statistical load balancing is maintained whenever a port in a LAG is lost or returned to service.

In Figure 13, four ports are aggregated together between the SI4091 and the enterprise routing device. Connectivity is maintained as long as one of the links remain active. The links to the server are also aggregated, allowing the secondary NIC to take over in the event that the primary NIC link fails.

Figure 13. Aggregating Ports for Link Redundancy

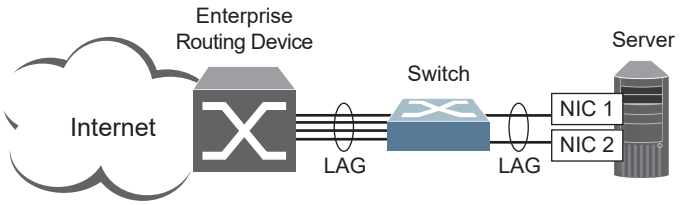

Default LAGs

By default, all available external ports are automatically assigned to an LACP dynamic LAG to allow plug-and-play operation.

Static LAGs

When you create and enable a static LAG, the LAG members (ports) take on certain settings necessary for correct operation of the aggregation feature.

Before Configuring Static LAGs

Before you configure your LAG, you must consider these settings, along with specific configuration rules, as follows:

- Read the configuration rules provided in the section, "Static LAG Configuration Rules" on page 145."
- Determine which SI4091 ports are to become *LAG members* (the specific ports making up the LAG).
- Ensure that the chosen ports are set to `enabled`.
- Ensure all member ports in a LAG have the same VLAN configuration.
- Consider how existing VLANs will be affected by the addition of a LAG.

Static LAG Configuration Rules

The aggregation feature operates according to specific configuration rules. When creating LAGs, consider the following rules that determine how a LAG reacts in any network topology:

- All LAGs must originate from one network entity and lead to one destination entity. For example, you cannot combine links from two different servers into one LAG.
- Any physical SI4091 port can belong to only one LAG.
- Internal ports (INT*x*) and external ports (EXT*x*) cannot become members of the same LAG.
- Depending on port availability, the SI4091 supports up to 14 internal ports in a LAG, or up to 10 external ports in a LAG.
- Aggregation from third-party devices must comply with Cisco® EtherChannel® technology.
- All LAG member ports must be assigned to the same VLAN configuration before the LAG can be enabled.
- You cannot change the VLAN settings of any subset of individual LAG members. VLAN settings must be configured for all LAG members at the same time.
- When an active port is configured in a LAG, the port becomes a *LAG member* when you enable the LAG.
- All ports in static LAGs must have the same link configuration (speed, duplex, flow control).
- When defining LAGs, if any of the desired ports is already a member of an existing LAG or LACP LAG, you must also remove those ports from their old LAG or LACP LAG.

Configuring a Static LAG

In the following example, three ports are aggregated between the SI4091 (noted as the *system switch*) and an external application switch.

Figure 14. LAG Group Configuration Example

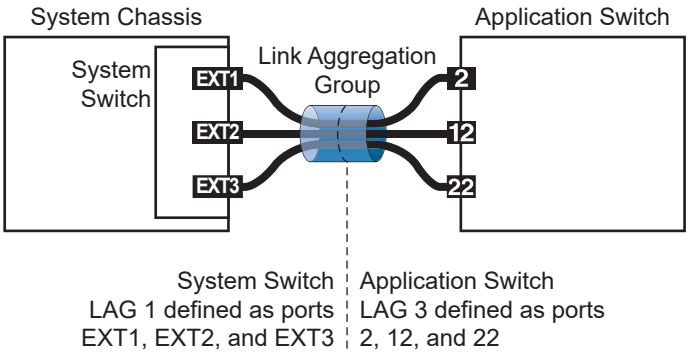

Before configuring each device in this example, you must connect to the appropriate device as the administrator.

1. Connect the ports that will be members in the LAG.
2. Configure the LAG using these steps on the SI4091:

   a. Define a LAG..

```
SI 4091(config)# portchannel 1 port ext1,ext2,ext3
SI 4091(config)# portchannel 1 enable
```

   b. Verify the configuration

```
SI 4091(config)# show portchannel information
```

   Examine the resulting information. If any settings are incorrect, make appropriate changes.

3. Repeat the process on the application switch.

```
RS G8011(config)# portchannel 3 port 2,12,22
RS G8011(config)# portchannel 3 enable
```

LAG 1 (on the SI4091) is now connected to LAG 3 on the application switch.

Note: In this example, a SI4091 and a Lenovo application switch are used. If a third-party device supporting link aggregation is used (such as Cisco routers and switches with EtherChannel technology or Sun's Quad Fast Ethernet Adapter), LAGs on the third-party device should be configured manually. Connection problems could arise when using automatic LAG negotiation on the third-party device.

4. Examine the aggregation information on each device. For example:.

```
SI 4091(config)# show portchannel information
PortChannel 1: Enabled
Protocol - Static
Port State:
  EXT1: forwarding
  EXT2: forwarding
  EXT3: forwarding
```

Information about each port in each configured LAG is displayed. Make sure that LAGs consist of the expected ports and that each port is in the expected state.

Configurable LAG Hash Algorithm

Traffic in a LAG is statistically distributed among member ports using a *hash* process where various address and attribute bits from each transmitted frame are recombined to specify the particular LAG port the frame will use. The SI4091 uses the RTAG7 model for LAG hashing.

The SI4091 can be configured to use a variety of hashing options. To achieve the most even traffic distribution, select options that exhibit a wide range of values for your particular network. Avoid hashing on information that is not usually present in the expected traffic, or which does not vary.

The SI4091 supports the following hashing options, which can be used in any combination:

- For Layer 2 traffic, one of the following must be applied:
  - Layer 2 source MAC address (`smac`)

```
SI 4091(config)# [no] portchannel thash l2thash
l2-source-mac-address
```

  - Layer 2 destination MAC address (`dmac`)

```
SI 4091(config)# [no] portchannel thash l2thash
l2-destination-mac-address
```

  - Both Layer 2 source and destination MAC address (the default)

```
SI 4091(config)# [no] portchannel thash l2thash
l2-source-destination-mac
```

Note: At least one Layer 2 option must always be enabled; The `smac` and `dmac` options may not both be disabled at the same time.

- For Layer 3 IPv4/IPv6 traffic, one of the following is permitted:
  - Layer 3 IPv4/IPv6 source IP address (`sip`)

```
SI 4091(config)# [no] portchannel thash l3thash l3-source-ip-address
```

  - Layer 3 IPv4/IPv6 destination IP address (`dip`)

```
SI 4091(config)# [no] portchannel thash l3thash
l3-destination-ip-address
```

  - Both Layer 3 source and destination IPv4/IPv6 address (the default)

```
SI 4091(config)# [no] portchannel thash l3thash
l3-source-destination-ip
```

  - Layer 2 hash override

    If Layer 2 hashing is preferred for Layer 3 traffic, disable the Layer 3 `sip` and `dip` hashing options and enable the `use-L2` option:

```
SI 4091(config)# [no] portchannel thash l3thash l3-use-l2-hash
```

    Layer 3 traffic will then use Layer 2 options for hashing.

- Layer 4 port information hash (disabled by default)

```
SI 4091(config)# [no] portchannel thash l4port
```

  When enabled, Layer 4 port information (TCP, UPD, etc.) is added to the hash if available. The `l4port` option is ignored when Layer 4 information is not included in the packet (such as for Layer 2 packets), or when the `use-L2` option is enabled.

Note: For MPLS packets, Layer 4 port information is excluded from the hash calculation. Instead, other IP fields are used, along with the first two MPLS labels.

- Ingress port number hash (disabled by default)

```
SI 4091(config)# [no] portchannel thash ingress
```

- For Fibre Channel over Ethernet (FCoE) traffic, any number of the following may be used:
  - FCoE destination ID (enabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe destination-id
```

  - FCoE source ID (enabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe source-id
```

  - FCoE CNTAG ID (disabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe cntag-id
```

  - FCoE fabric ID (disabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe fabric-id
```

  - FCoE originator transmit ID (disabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe originator-id
```

  - FCoE responder transmit ID (disabled by default)

```
SI 4091(config)# [no] portchannel thash fcoe responder-id
```

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic LAG or Link Aggregation Group) with any device that supports the standard. Please refer to IEEE 802.3ad-2002 for a full description of the standard.

LACP Overview

IEEE 802.3ad allows standard Ethernet links to form a single Layer 2 link using the Link Aggregation Control Protocol (LACP). Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex, and treating them as if they were part of a single, logical link segment. If a link in a LACP LAG fails, traffic is reassigned dynamically to the remaining link or links of the dynamic LAG.

In general, LACP must be configured on both the SI4091 as well as on the connecting device. This section describes LACP on the SI4091. For configuring LACP on the connecting device, see the appropriate product documentation.

The SI4091 supports up to 52 LACP LAGs, each with up to 24 ports.

Note: LACP implementation in Enterprise NOS does not support the Churn machine, an option used to detect if the port is operable within a bounded time period between the actor and the partner. Only the Marker Responder is implemented, and there is no marker protocol generator.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

- **System ID**: an integer value based on the SI4091's MAC address and the system priority assigned in the CLI.
- **Admin key:** a port's *admin key* is an integer value (1 - 65535) that you can configure in the CLI. Each SI4091 port that participates in the same LACP LAG must have the same *admin key* value. The admin key is *local significant*, which means the partner device does not need to use the same admin key value.

For example, consider two devices, the SI4091 and a Partner switch, as shown in Table 14.

Table 14. *Actor vs. Partner LACP configuration*

SI4091	Partner Switch
Port EXT1 (LAG = 25, admin key = 100)	Port 1 (admin key = 50)
Port EXT2 (LAG = 25, admin key = 100)	Port 2 (admin key = 50)
Port EXT3 (LAG = 26, admin key = 100)	Port 3 (admin key = 70)

In the configuration shown in Table 14, the SI4091 ports EXT1 and EXT2 are aggregated to form an LACP LAG with Partner switch ports 1 and 2. Only ports with the same LAG ID are aggregated in the LAG. The SI4091 port EXT3 is not aggregated in the LAG because it has a different LAG ID. Other ports configured with the same admin key on the SI4091 but with a different LAG ID (due to Partner

switch admin key configuration or due to Partner switch MAC address being different) can be aggregated in another LAG. For instance, SI4091 port EXT3 can be aggregated in another LAG with ports that have the same LAG ID as port EXT3.

To prevent SI4091 ports (with the same admin key) from aggregating in another LAG, you can configure a LAG ID. Ports with the same admin key (although with different LAG IDs) compete to get aggregated into a LAG. The LAG ID for the LAG is decided based on the first port that is aggregated in the group. Ports with this LAG ID get aggregated and the other ports are placed in `suspended` mode. As per the configuration shown in Table 14 on page 151, if port EXT1 gets aggregated first, then the LAG ID of port EXT1 would be the LAG ID of the LAG. Port EXT3 would be placed in suspended mode. When in suspended mode, a port transmits only LACP data units (LACPDUs) and discards all other traffic.

A port may also be placed in suspended mode for the following reasons:

- When LACP is configured on the port but it stops receiving LACPDUs from the partner switch.
- When the port has a different LAG ID because of the partner switch MAC being different. For example: when the SI4091 is connected to two partners.

The LAG ID can be viewed using the following command:

```
SI 4091 # show lacp information
```

LACP provides for the controlled addition and removal of physical links for the link aggregation.

Each port in the SI4091 can have one of the following LACP modes.

- `off` (default)

  The user can configure this port in to a regular static LAG.

- `active`

  The port is capable of forming an LACP LAG. This port sends LACPDU packets to partner system ports.

- `passive`

  The port is capable of forming an LACP LAG. This port only responds to the LACPDU packets sent from an LACP *active* port.

Each active LACP port transmits LACP data units (LACPDUs), while each passive LACP port listens for LACPDUs. During LACP negotiation, the admin key is exchanged. The LACP LAG is enabled as long as the information matches at both ends of the link. If the admin key value changes for a port at either end of the link, that port's association with the LACP LAG is lost.

When the system is initialized, all ports by default are in LACP *off* mode and are assigned unique *admin keys*. To make a group of ports available for aggregation, you assign them all the same *admin key*. You must set the port's LACP mode to *active* to activate LACP negotiation. You can set other port's LACP mode to passive, to reduce the amount of LACPDU traffic at the initial LAG-forming stage.

Note: When the LACPDU rate is lower than the control plane policing (CoPP) LACP queue, LACP does not display logging messages because the configurable queue has exceeded the CPU capacity.

In the information displays, static LAGs are listed as LAGs 1 though 52. Dynamic LACP LAGs are listed as 25 through 48.

Configuring LACP

Use the following procedure to configure LACP for port EXT1 and port EXT2 to participate in link aggregation.

1. Configure port parameters. All ports that participate in the LACP LAG must have the same settings, including VLAN membership.
2. Define a portchannel by assigning an ID and an LACP admin key:

```
SI 4091(config)# portchannel 25 lacp key 100
```

3. Select the ports and bind the LACP admin key. Only ports with the same admin key can form an LACP LAG.

```
SI 4091(config)# interface port ext1-ext2
SI 4091(config-if)# lacp key 100
```

Note: If any of the desired ports is a member of an existing LAG, you must also remove the desired ports from their old LAG.

4. Set the LACP mode.

```
SI 4091(config-if)# lacp mode active
SI 4091(config-if)# exit
```

LACP Individual

Ports assigned with an LACP admin key are prevented by default from forming individual links if they cannot join an LACP LAG. To override this behavior, use the following commands:

```
SI 4091(config)# interface port <port no or range>
SI 4091(config-if)# no lacp suspend-individual
```

This allows the selected ports to be treated as normal link-up ports if they do not receive any LACPDUs.

To configure the LACP individual setting for all the ports in an LACP LAG, use the following commands:

```
SI 4091(config-if)# interface portchannel lacp <adminkey of the LAG>
SI 4091(config-if)# [no] lacp suspend-individual
```

By default, external ports are configured with `lacp suspend-individual`, while internal ports are configured with `no lacp suspend-individual`.

When you use the command `no lacp suspend-individual` to configure uplink ports to forward traffic in case LACPDUs are no longer received, *one* LACP member port will change to LACP individual mode, and only that single LACP port will forward traffic.

Chapter 11. Quality of Service

Quality of Service (QoS) features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate QoS level.

The following topics are discussed in this section:

QoS Overview

QoS helps you allocate guaranteed bandwidth to critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Also, you can put a high priority on applications that are sensitive to timing out or those that cannot tolerate delay, assigning that traffic to a high-priority queue.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

Figure 15 on page 156 shows the basic QoS model used by the SI4091 System Interconnect Module (SI4091).

Figure 15. QoS Model

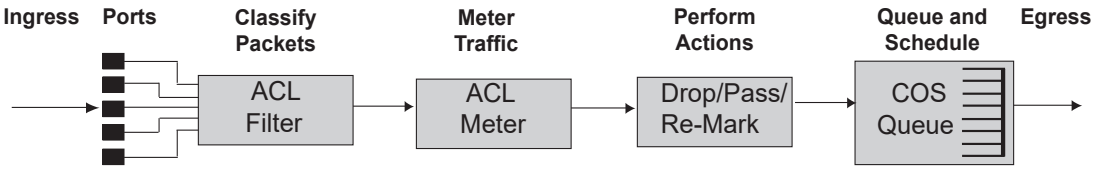

The SI4091 uses the Differentiated Services (DiffServ) architecture to provide QoS functions. DiffServ is described in IETF RFC 2474 and RFC 2475.

With DiffServ, you can establish policies for directing traffic. A policy is a traffic-controlling mechanism that monitors the characteristics of the traffic (for example, its source, destination, and protocol) and performs a controlling action on the traffic when certain characteristics are matched.

The SI4091 can classify traffic by reading the DiffServ Code Point (DSCP) or IEEE 802.1p priority value, or by using filters to match specific criteria. When network traffic attributes match those specified in a traffic pattern, the policy instructs the SI4091 to perform specified actions on each packet that passes through it. The packets are assigned to different Class of Service (COS) queues and scheduled for transmission.

The basic SI4091 QoS model works as follows:

- Classify traffic:
  - Read DSCP
  - Read 802.1p Priority
  - Match ACL filter parameters
- Meter traffic:
  - Define bandwidth and burst parameters
  - Select actions to perform on in-profile and out-of-profile traffic

- Perform actions:
  - Drop packets
  - Pass packets
  - Mark DSCP or 802.1p Priority
  - Set COS queue (with or without re-marking)
- Queue and schedule traffic:
  - Place packets in one of the available COS queues
  - Schedule transmission based on the COS queue weight

Using ACL Filters

Access Control Lists (ACLs) are filters that allow you to classify and segment traffic, so you can provide different levels of service to different traffic types. Each filter defines conditions that packets must match for inclusion in a particular service class, and also the actions that are performed for matching traffic.

The SI4091 allows you to classify packets based on various parameters. For example:

- Ethernet—source MAC, destination MAC, VLAN number/mask, Ethernet type, priority
- IPv4—source IP address/mask, destination address/mask, type of service, IP protocol number
- IPv6—source IP address/prefix, destination address/prefix, next header, flow label, traffic class
- TCP/UPD—source port, destination port, TCP flag
- Packet format—Ethernet format, tagging format, IPv4, IPv6
- Egress port

For ACL details, see "Access Control Lists" on page 101.

Summary of ACL Actions

Actions determine how the traffic is treated. The SI4091 QoS actions include the following:

- Pass or Drop the packet
- Re-mark the packet with a new DiffServ Code Point (DSCP)
- Re-mark the 802.1p field
- Set the COS queue

ACL Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the SI4091 by configuring a QoS meter (if desired) and assigning ACL Groups to ports. When you add ACL Groups to a port, make sure they are ordered correctly in terms of precedence.

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- **In-Profile**–If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile**–If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Note: Metering is not supported for IPv6 ACLs. All traffic matching an IPv6 ACL is considered in-profile for re-marking purposes.

Using meters, you set a Committed Rate in Kbps (1000 bits per second in each Kbps). All traffic within this Committed Rate is In-Profile. Additionally, you can set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level traffic should receive.
- Change the 802.1p priority of a packet.

Using DSCP Values to Provide QoS

The six most significant bits in the TOS byte of the IP header are defined as DiffServ Code Points (DSCP). Packets are marked with a certain value depending on the type of treatment the packet must receive in the network device. DSCP is a measure of the Quality of Service (QoS) level of the packet.

Differentiated Services Concepts

To differentiate between traffic flows, packets can be classified by their DSCP value. The Differentiated Services (DS) field in the IP header is an octet, and the first six bits, called the DS Code Point (DSCP), can provide QoS functions. Each packet carries its own QoS state in the DSCP. There are 64 possible DSCP values (0-63).

Figure 16. Layer 3 IPv4 Packet

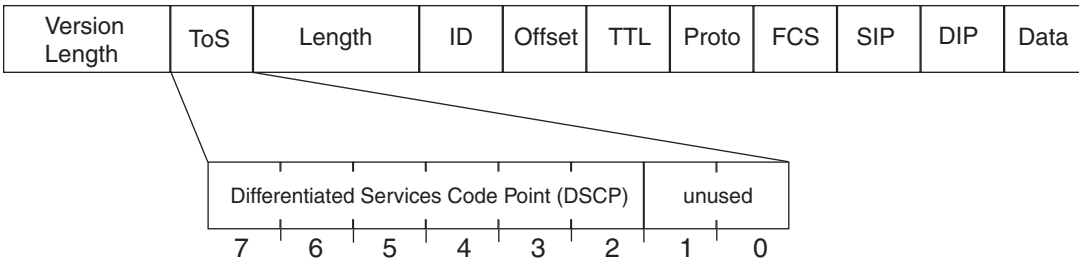

The SI4091 can perform the following actions to the DSCP:

- Read the DSCP value of ingress packets
- Re-mark the DSCP value to a new value
- Map the DSCP value to an 802.1p priority

Once the DSCP value is marked, the SI4091 can use it to direct traffic prioritization.

Per-Hop Behavior

The DSCP value determines the Per Hop Behavior (PHB) of each packet. The PHB is the forwarding treatment given to packets at each hop. QoS policies are built by applying a set of rules to packets, based on the DSCP value, as they hop through the network.

The SI4091 default settings are based on the following standard PHBs, as defined in the IEEE standards:

- Expedited Forwarding (EF)—This PHB has the highest egress priority and lowest drop precedence level. EF traffic is forwarded ahead of all other traffic. EF PHB is described in RFC 2598.
- Assured Forwarding (AF)—This PHB contains four service levels, each with a different drop precedence, as shown below. Routers use drop precedence to determine which packets to discard last when the network becomes congested. AF PHB is described in RFC 2597.

Drop Precedence	Class 1	Class 2	Class 3	Class 4
Low	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Medium	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High	AF13 (DSCP 14)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

- Class Selector (CS)—This PHB has eight priority classes, with CS7 representing the highest priority, and CS0 representing the lowest priority, as shown below. CS PHB is described in RFC 2474.

Priority	Class Selector	DSCP
Highest	CS7	56
	CS6	48
	CS5	40
	CS4	32
	CS3	24
	CS2	16
	CS1	8
Lowest	CS0	0

QoS Levels

Table 15 shows the default service levels provided by the SI4091, listed from highest to lowest importance:

Table 15. *Default QoS Service Levels*

Service Level	Default PHB	802.1p Priority
Critical	CS7	7
Network Control	CS6	6
Premium	EF, CS5	5
Platinum	AF41, AF42, AF43, CS4	4
Gold	AF31, AF32, AF33, CS3	3
Silver	AF21, AF22, AF23, CS2	2
Bronze	AF11, AF12, AF13, CS1	1
Standard	DF, CS0	0

DSCP Re-Marking and Mapping

The SI4091 can re-mark the DSCP value of ingress packets to a new value, and set the 802.1p priority value, based on the DSCP value. You can view the settings by using the following command:

```
SI 4091# show qos dscp
Current DSCP Remarking Configuration: OFF

  DSCP    New DSCP  New 802.1p Prio
--------  --------  ---------------
    0         0          0
    1         1          0
...
   51        51          0
   52        52          0
   53        53          0
   54        54          0
   55        55          0
   56        56          7
   57        57          0
   58        58          0
   59        59          0
   60        60          0
   61        61          0
   62        62          0
   63        63          0
```

Use the following command to turn on DSCP re-marking globally:

```
SI 4091(config)# qos dscp re-marking
```

You must then enable DSCP re-marking on any port that you wish to perform this function.

Note: If an ACL meter is configured for DSCP re-marking, the meter function takes precedence over QoS re-marking.

DSCP Re-Marking Configuration Example

1. Turn DSCP re-marking on globally, and define the DSCP-DSCP-802.1p mapping. You can use the default mapping.

```
SI 4091(config)# qos dscp re-marking
SI 4091(config)# qos dscp dscp-mapping <DSCP value (0-63)> <new value>
SI 4091(config)# qos dscp dot1p-mapping <DSCP value (0-63)> <802.1p value>
```

2. Enable DSCP re-marking on a port.

```
SI 4091(config)# interface port ext1
SI 4091(config-if)# qos dscp re-marking
SI 4091(config-if)# exit
```

Using 802.1p Priorities to Provide QoS

Enterprise NOS provides Quality of Service functions based on the priority bits in a packet's VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1q VLAN header.) The 802.1p bits, if present in the packet, specify the priority that should be given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority bit value.

The IEEE 802.1p standard uses eight levels of priority (0-7). Priority 7 is assigned to highest priority network traffic, such RIP routing table updates, priorities 5-6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a "best effort" traffic prioritization, and this is the default when traffic priority has not been configured on your network. The SI4091 can filter packets based on the 802.1p values, and it can assign or overwrite the 802.1p value in the packet.

Figure 17. Layer 2 802.1q/802.1p VLAN Tagged Packet

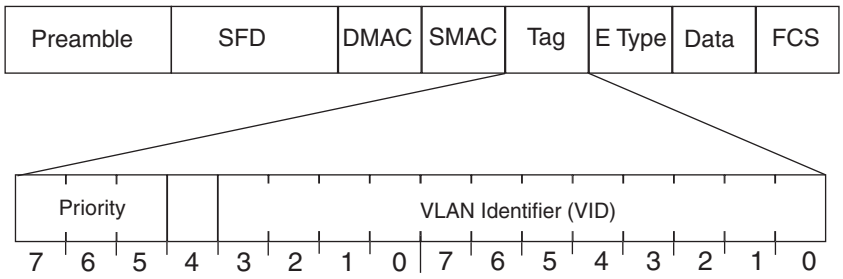

Ingress packets receive a priority value, as follows:

- **Tagged packets**—SI4091 reads the 802.1p priority in the VLAN tag.
- **Untagged packets**—SI4091 tags the packet and assigns an 802.1p priority, based on the port's default priority.

Egress packets are placed in a COS queue based on the priority value, and scheduled for transmission based on the scheduling weight of the COS queue.

To configure a port's default 802.1p priority value, use the following commands.

```
SI 4091(config)# interface port ext1
SI 4091(config-if)# dot1p <802.1p value (0-7)>
SI 4091(config-if)# exit
```

Queuing and Scheduling

The SI4091 can be configured to have either 2 or 8 output Class of Service (COS) queues per port, into which each packet is placed. Each packet's 802.1p priority determines its COS queue, except when an ACL action sets the COS queue of the packet.

You can configure the following attributes for COS queues:

- Map 802.1p priority value to a COS queue
- Define the scheduling weight of each COS queue

You can map 802.1p priority value to a COS queue, as follows:

```
SI 4091(config)# qos transmit-queue mapping <802.1p priority value (0-7)> <COS queue
    (0-7)>
```

To set the COS queue scheduling weight, use the following command.

```
SI 4091(config)# qos transmit-queue weight-cos <COSq number>
    <COSq weight (0-15)>
```

The scheduling weight can be set from 0 to 15. Weight values from 1 to 15 set the queue to use weighted round-robin (WRR) scheduling, which distributes larger numbers of packets to queues with the highest weight values. For distribution purposes, each packet is counted the same, regardless of the packet's size.

A scheduling weight of 0 (zero) indicates strict priority. Traffic in strict priority queue has precedence over other all queues. If more than one queue is assigned a weight of 0, the strict queue with highest queue number will be served first. Once all traffic in strict queues is delivered, any remaining bandwidth will be allocated to the WRR queues, divided according to their weight values.

Note: Use caution when assigning strict scheduling to queues. Heavy traffic in queues assigned with a weight of 0 can starve lower priority queues.

For a scheduling method that uses a weighted deficit round-robin (WDRR) algorithm, distributing packets with an awareness of packet size, see "Enhanced Transmission Selection" on page 213.

Control Plane Protection

The switch's *control plane* receives and processes packets required for the stable function of its various traffic protocols. Control plane traffic usually occurs at a low rate. However, in some situations, such as during Denial Of Service (DOS) attacks, the switch may receive this traffic at a significantly higher rate. If the control plane is unable to process the received control traffic, the switch operation may become unstable.

The control plane receives packets channeled through protocol-specific packet queues. Multiple protocols can be channeled through a common packet queue, but any specific protocol cannot be channeled through multiple queues. These packet queues are applicable only to the packets received by the software and does not impact the regular switching or routing traffic. Packet queues with a higher number have higher priority.

You can configure the bandwidth for each packet queue. Protocols that share a packet queue will also share the bandwidth.

Use the following command to assign a protocol to a specific queue:

```
SI 4091(config)# qos protocol-packet-control packet-queue-map <queue (0-44)>
     <protocol>
```

Use the following command to specify a rate limit, in packets per second, to a specific queue:

```
SI 4091(config)# qos protocol-packet-control rate-limit-packet-queue
      <queue (0-44)>  <rate (1-10000)>
```

Packet Drop Logging

Packet drop logging allows you to monitor network deficiencies by generating system log messages for packets dropped by the switch control plane queues. By default, the switch will generate packet drop logging messages once every two minutes, specifying the type of traffic, queue data rate and queue number on which the drops occurred. For example:

```
Apr 19 11:27:35 172.31.37.200 NOTICE Protocol control discards: ARP
Broadcast
packets are received at rate higher than 200pps, hence are discarded on
queue 5.
```

To enable or disable packet drop logging, use the following commands:

```
SI 4091(config)# [no] logging pdrop enable
```

You can adjust the logging interval from 0 to 30 minutes using the following command:

```
SI 4091(config)# logging pdrop interval <minutes (0-30)>
```

Setting the logging interval to 0 will cause a packet drop logging message to be generated immediately (with up to a one second delay), and will then ignore further drops on the same queue during the next two minutes.

Setting the logging interval to a greater value (from 1 to 30 minutes) will cause packet drop logging messages to be generated at the specified regular time intervals. Once the packet drops stop, or if new packet drops are encountered only within two minutes after a system log message, the switch will not display additional messages.

Note: If you need a packet drop interval of 30 minutes, you must set it manually using the command:

```
SI 4091(config)# logging pdrop interval 30
```

Part 4: IP Routing

This section discusses Layer 3 switching functions. In addition to switching traffic at near line rates, the application switch can perform multi-protocol routing. This section discusses basic routing and advanced routing protocols:

- Basic Routing
- IPv6 Host Management
- Internet Group Management Protocol (IGMP)

Chapter 12. Basic IP Addresses

This chapter provides configuration background and examples for using the SI4091 System Interconnect Module (SI4091) to perform BOOTP and DHCP functions.

Dynamic Host Configuration Protocol (DHCP) provides a framework for automatically assigning IP addresses and configuration information to other IP hosts or clients in a large TCP/IP network. Without DHCP, the IP address must be entered manually for each network device. DHCP allows a network administrator to distribute IP addresses from a central point and automatically send a new IP address when a device is connected to a different place in the network.

The switch accepts gateway configuration parameters if they have not been configured manually. The switch ignores DHCP gateway parameters if the gateway is configured.

DHCP is an extension of another network IP management protocol, Bootstrap Protocol (BOOTP), with an additional capability of being able to dynamically allocate reusable network addresses and configuration parameters for client operation.

Built on the client/server model, DHCP allows hosts or clients on an IP network to obtain their configurations from a DHCP server, thereby reducing network administration. The most significant configuration the client receives from the server is its required IP address; (other optional parameters include the “generic” file name to be booted, the address of the default gateway, and so forth).

To enable the DHCP client on the SI4091, use the following command:

```
SI 4091(config)# system dhcp
```

Chapter 13. Internet Protocol Version 6

Internet Protocol version 6 (IPv6) is a network layer protocol intended to expand the network address space. IPv6 is a robust and expandable protocol that meets the need for increased physical address space. The switch supports the following RFCs for IPv6-related features:

- RFC 1981
- RFC 2451
- RFC 2460
- RFC 2461
- RFC 2462
- RFC 2474
- RFC 2526
- RFC 2711
- RFC 3289
- RFC 3306
- RFC 3307
- RFC 3411
- RFC 3412
- RFC 3413
- RFC 3414
- RFC 3484
- RFC 3602
- RFC 3879
- RFC 4007
- RFC 4213
- RFC 4291
- RFC 4293
- RFC 4293
- RFC 4443
- RFC 4861
- RFC 4862
- RFC 5095

This chapter describes the basic configuration of IPv6 addresses and how to manage the switch via IPv6 host management.

IPv6 Limitations

The following IPv6 feature restrictions apply to this release.

- Dynamic Host Control Protocol for IPv6 (DHCPv6) is supported only on the management interface (MGT-port).
- Routing Information Protocol for IPv6 (RIPng) is not supported.

Most other Enterprise NOS 8.4 features permit IP addresses to be configured using either IPv4 or IPv6 address formats. However, the following switch features support IPv4 only:

- Default switch management IP address
- Bootstrap Protocol (BOOTP)
- DHCP for internal management interfaces 127 and 128
- RADIUS, TACACS+ and LDAP
- QoS metering and re-marking ACLs for out-profile traffic

IPv6 Address Format

The IPv6 address is 128 bits (16 bytes) long and is represented as a sequence of eight 16-bit hex values, separated by colons.

Each IPv6 address has two parts:

- Subnet prefix representing the network to which the interface is connected
- Local identifier, either derived from the MAC address or user-configured

The preferred hexadecimal format is as follows:

```
xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx
```

Example IPv6 address:

```
FEDC:BA98:7654:BA98:FEDC:1234:ABCD:5412
```

Some addresses can contain long sequences of zeros. A single contiguous sequence of zeros can be compressed to :: (two colons). For example, consider the following IPv6 address:

```
FE80:0:0:0:2AA:FF:FA:4CA2
```

The address can be compressed as follows:

```
FE80::2AA:FF:FA:4CA2
```

Unlike IPv4, a subnet mask is not used for IPv6 addresses. IPv6 uses the subnet prefix as the network identifier. The prefix is the part of the address that indicates the bits that have fixed values or are the bits of the subnet prefix. An IPv6 prefix is written in address/prefix-length notation. For example, in the following address, 64 is the network prefix:

```
21DA:D300:0000:2F3C::/64
```

IPv6 addresses can be either user-configured or automatically configured. Automatically configured addresses always have a 64-bit subnet prefix and a 64-bit interface identifier. In most implementations, the interface identifier is derived from the switch's MAC address, using a method called EUI-64.

Most Enterprise NOS 8.4 features permit IP addresses to be configured using either IPv4 or IPv6 address formats. Throughout this manual, *IP address* is used in places where either an IPv4 or IPv6 address is allowed. In places where only one type of address is allowed, the type (*IPv4* or *IPv6* is specified).

IPv6 Address Types

IPv6 supports three types of addresses: unicast (one-to-one), multicast (one-to-many), and anycast (one-to-nearest). Multicast addresses replace the use of broadcast addresses.

Unicast Address

Unicast is a communication between a single host and a single receiver. Packets sent to a unicast address are delivered to the interface identified by that address. IPv6 defines the following types of unicast address:

- Global Unicast address: An address that can be reached and identified globally. Global Unicast addresses use the high-order bit range up to FF00, therefore all non-multicast and non-link-local addresses are considered to be global unicast. A manually configured IPv6 address must be fully specified. Autoconfigured IPv6 addresses are comprised of a prefix combined with the 64-bit EUI. RFC 4291 defines the IPv6 addressing architecture.

  The interface ID must be unique within the same subnet.

- Link-local unicast address: An address used to communicate with a neighbor on the same link. Link-local addresses use the format `FE80::EUI`

  Link-local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or when no routers are present.

  Routers must not forward any packets with link-local source or destination addresses to other links.

Multicast

Multicast is communication between a single host and multiple receivers. Packets are sent to all interfaces identified by that address. An interface may belong to any number of multicast groups.

A multicast address (FF00 - FFFF) is an identifier for a group interface. The multicast address most often encountered is a solicited-node multicast address using prefix `FF02::1:FF00:0000/104` with the low-order 24 bits of the unicast or anycast address.

The following well-known multicast addresses are pre-defined. The group IDs defined in this section are defined for explicit scope values, as follows:

`FF00:::::::0` through `FF0F:::::::0`

Anycast

Packets sent to an anycast address or list of addresses are delivered to the nearest interface identified by that address. Anycast is a communication between a single sender and a list of addresses.

Anycast addresses are allocated from the unicast address space, using any of the defined unicast address formats. Thus, anycast addresses are syntactically indistinguishable from unicast addresses. When a unicast address is assigned to

more than one interface, thus turning it into an anycast address, the nodes to which the address is assigned must be explicitly configured to know that it is an anycast address.

IPv6 Address Autoconfiguration

IPv6 supports the following types of address autoconfiguration:

- **Stateful address configuration**

  Address configuration is based on the use of a stateful address configuration protocol, such as DHCPv6, to obtain addresses and other configuration options.

- **Stateless address configuration**

  Address configuration is based on the receipt of Router Advertisement messages that contain one or more Prefix Information options.

Enterprise NOS 8.4 supports stateless address configuration.

Stateless address configuration allows hosts on a link to configure themselves with link-local addresses and with addresses derived from prefixes advertised by local routers. Even if no router is present, hosts on the same link can configure themselves with link-local addresses and communicate without manual configuration.

IPv6 Management Interfaces

Management interfaces 125 and 126 support multiple IPv6 addresses. You can manually configure up to two IPv6 addresses for each management interface, or you can allow the switch to use stateless autoconfiguration. By default, the switch automatically configures the IPv6 address of its management interfaces.

You can manually configure two IPv6 addresses for each interface, as follows:

- Initial IPv6 address is a global unicast or anycast address:

  interface ip *<x>*
  address *<IPv6 address>*

  Note that you cannot configure both addresses as anycast. If you configure an anycast address on the interface you must also configure a global unicast address on that interface.

- Second IPv6 address can be a unicast or anycast address:

  interface ip *<x>*
  secaddr6 *<IPv6 address>*

You cannot configure an IPv4 address on an IPv6 management interface. Each interface can be configured with only one address type: either IPv4 or IPv6, but not both.

The SI4091 supports only management IP addresses; any IP address configured will automatically belong to the management VLAN (4095).

Interface 125/126 is reserved for IPv6 host support. This interface is included in management VLAN 4095. Interface 125 configures the IP address for the EXTM port and interface 126 configures the IP address for the MGT1 port. Use IPv6 gateway mode to configure the IPv6 gateways (**ip gateway6** {*<gateway number>*} **address** *<IPv6 address>*).

IPv6 gateway 3 and 4 are the default IPv6 management gateways.

Neighbor Discovery

The switch uses Neighbor Discovery protocol (ND) to gather information about other router and host nodes, including the IPv6 addresses. Host nodes use ND to configure their interfaces and perform health detection. ND allows each node to determine the link-layer addresses of neighboring nodes, and to keep track of each neighbor's information. A neighboring node is a host or a router that is linked directly to the switch. The switch supports Neighbor Discovery as described in RFC 4861.

Neighbor Discover messages allow network nodes to exchange information, as follows:

- *Neighbor Solicitations* allow a node to discover information about other nodes.
- *Neighbor Advertisements* are sent in response to Neighbor Solicitations. The Neighbor Advertisement contains information required by nodes to determine the link-layer address of the sender, and the sender's role on the network.
- IPv6 hosts use *Router Solicitations* to discover IPv6 routers. When a router receives a Router Solicitation, it responds immediately to the host.
- Routers uses *Router Advertisements* to announce its presence on the network, and to provide its address prefix to neighbor devices. IPv6 hosts listen for Router Advertisements, and uses the information to build a list of default routers. Each host uses this information to perform autoconfiguration of IPv6 addresses.
- *Redirect messages* are sent by IPv6 routers to inform hosts of a better first-hop address for a specific destination. Redirect messages are only sent by routers for unicast traffic, are only unicast to originating hosts, and are only processed by hosts.

ND configuration for various advertisements, flags, and interval settings is performed on a per-interface basis using the following commands:

```
SI 4091(config)# interface ip <interface number>
SI 4091(config-ip-if)# ipv6 nd <command>
```

Supported Applications

The following applications have been enhanced to provide IPv6 support.

- **Ping**

  The **ping** command supports IPv6 addresses. Use the following format to ping an IPv6 address:

  ping *<host name>*|*<IPv6 address>* [**-n** *<tries (0-4294967295)>*] [**-w** *<msec delay (0-4294967295)>*] [**-l** *<length (0/32-65500/2080)>*] [**-s** *<IP source>*] [**-v** *<TOS (0-255)>*] [**-f**] [**-t**]

  To ping a link-local address (begins with FE80), provide an interface index, as follows:

  ping *<IPv6 address>***%***<Interface index>* [**-n** *<tries (0-4294967295)>*] [**-w** *<msec delay (0-4294967295)>*] [**-l** *<length (0/32-65500/2080)>*] [**-s** *<IP source>*] [**-v** *<TOS (0-255)>*] [**-f**] [**-t**]

- **Traceroute**

  The **traceroute** command supports IPv6 addresses (but not link-local addresses). Use the following format to perform a traceroute to an IPv6 address:

  traceroute *<host name>*| *<IPv6 address>* [*<max-hops (1-32)>* [*<msec delay (1-4294967295)>*]]

- **Telnet**

  The **telnet** command supports IPv6 addresses, but not link-local addresses. Use the following format to Telnet to an IPv6 address:

  telnet *<host name>*| *<IPv6 address>* [*<port>*]

- **SSH**

  Secure Shell (SSH) connections over IPv6 are supported, but not link-local addresses. The following syntax is required from the client:

  ssh -u *<IPv6 address>*

  Example:

  ```
  ssh -u 2001:2:3:4:0:0:0:142
  ```

- **TFTP**

  The TFTP commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported.

- **FTP**

  The FTP commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported.

- **DNS client**

  DNS commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported. Use the following command to specify the type of DNS query to be sent first:

```
SI 4091(config)# ip dns ipv6 request-version {ipv4|ipv6}
```

  If you set the request version to `v4`, the DNS application sends an *A* query first, to resolve the hostname with an IPv4 address. If no *A* record is found for that hostname (no IPv4 address for that hostname) an *AAAA* query is sent to resolve the hostname with a IPv6 address.

  If you set the request version to `v6`, the DNS application sends an *AAAA* query first, to resolve the hostname with an IPv6 address. If no *AAAA* record is found for that hostname (no IPv6 address for that hostname) an *A* query is sent to resolve the hostname with an IPv4 address.

Configuration Guidelines

When you configure an interface for IPv6, consider the following guidelines:

- Support for subnet router anycast addresses is not available.
- Interfaces 125 and 126 are reserved for IPv6 management.
- A single interface can accept either IPv4 or IPv6 addresses, but not both IPv4 and IPv6 addresses.
- A single interface can accept multiple IPv6 addresses.
- A single interface can accept only one IPv4 address.
- If you change the IPv6 address of a configured interface to an IPv4 address, all IPv6 settings are deleted.
- Health checks are not supported for IPv6 gateways.
- IPv6 interfaces support Path MTU Discovery. The CPU's MTU is fixed at 1500 bytes.
- Support for jumbo frames (1,500 to 9,216 byte MTUs) is limited. Any jumbo frames intended for the CPU must be fragmented by the remote node. The switch can re-assemble fragmented packets up to 9k. It can also fragment and transmit jumbo packets received from higher layers.

IPv6 Configuration Examples

This section provides steps to configure IPv6 on the switch.

IPv6 Example 1

The following example uses IPv6 host mode to autoconfigure an IPv6 address for the management interface.

1. Enable IPv6 host mode on an interface.

```
SI 4091(config)# interface ip 125        (Select IP interface 125)
SI 4091(config-ip-if)# ip6host           (Enable IPv6 host mode
SI 4091(config-ip-if)# enable            (Enable the IP interface)
```

2. Configure (select and enable) the IPv6 default gateway.

```
SI 4091(config)# ip gateway6 3 address
2001:BA98:7654:BA98:FEDC:1234:ABCD:5412  enable
```

3. Verify the interface address.

```
SI 4091(config)# show interface ip 125(Display interface information)
```

IPv6 Example 2

Use the following example to manually configure IPv6 on an interface.

1. Assign an IPv6 address and prefix length to the interface.

```
SI 4091(config)# interface ip 125
SI 4091(config-ip-if)# ipv6 address
2001:BA98:7654:BA98:FEDC:1234:ABCD:5214 32 ena
SI 4091(config-ip-if)# ipv6 prefixlen 64
SI 4091(config-ip-if)# ipv6 secaddr6 address 2003::1 32
```

The secondary IPv6 address is compressed, and the prefix length is 32.

2. Configure the IPv6 default gateway.

```
SI 4091(config)# ip gateway6 3 address
2001:BA98:7654:BA98:FEDC:1234:ABCD:5412 ena(Enable default gateway)
```

3. Configure Router advertisements for the interface (optional)

```
SI 4091(config-ip-if)# ipv6 nd advmtu  (Enable Router Advertisements)
```

4. Verify the configuration.

```
SI 4091# show ipv6 interface            (View current IP settings)
SI 4091# show ipv6 gateway6 3           (View current IP gateway 3 settings)
```

Chapter 14. Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is used by IPv4 Multicast routers to learn about the existence of host group members on their directly attached subnet (see RFC 2236).

IPv4 Multicast routers acquire IGMP information by broadcasting IGMP Membership Queries and listening for IPv4 hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IPv4 Multicast source that provides the data streams and the clients that want to receive the data.

The SI4091 can perform IGMP Snooping.

Note: Enterprise NOS 8.4 does not support IPv6 for IGMP.

The following topics are discussed in this chapter:

IGMP Snooping

IGMP Snooping prevents multicast traffic from being flooded to all ports. Instead, the SI4091 learns which server hosts are interested in receiving multicast traffic, and forwards it only to ports connected to those servers. In this way, other ports are not burdened with unwanted multicast traffic.

The SI4091 detects IGMP Membership Reports from attached clients and acts as a proxy to set up a dedicated path between the requesting host and a local IPv4 Multicast router. After the pathway is established, the SI4091 blocks the IPv4 Multicast stream from flowing through any port that does not connect to a host member, thus conserving bandwidth.

The client-server path is set up as follows:

- An IPv4 Multicast Router (Mrouter) sends *Membership Queries* to the SI4091, which forwards them to all ports in a given VLAN.
- Hosts that want to receive the multicast data stream send *Membership Reports* to the SI4091, which sends a proxy Membership Report to the Mrouter.
- The SI4091 sets up a path between the Mrouter and the host, and blocks all other ports from receiving the multicast.
- Periodically, the Mrouter sends Membership Queries to ensure that the host wants to continue receiving the multicast. If a host fails to respond with a Membership Report, the Mrouter stops sending the multicast to that path.
- The host can send an IGMP Leave packet to the SI4091, which responds with an IGMP Groups Specific Query in order to check if there are other clients that want to receive the multicast traffic for the group referenced in the Leave packet. If an IGMP Report is not received, the group is deleted from the port and the multicast path is terminated. The SI4091 then sends a Proxy Leave packet to the Mrouter in order to updated it. If the FastLeave option is enabled on a VLAN, the multicast path is terminated immediately and the Leave packet is directly forwarded to the Mrouter.

IGMP Groups

The SI4091 supports a maximum of 256 IGMP entries, on a maximum of 1024 VLANs. One IGMP entry is allocated for each unique join request, based on the VLAN and IGMP group address only (regardless of the port). If multiple ports join the same IGMP group using the same VLAN, only a single IGMP entry is used.

IGMPv3

IGMPv3 includes new membership report messages to extend IGMP functionality. The SI4091 provides snooping capability for all types of IGMP version 3 (IGMPv3) Membership Reports, as described in RFC 3376.

IGMPv3 supports Source-Specific Multicast (SSM). SSM identifies session traffic by both source and group addresses. The SI4091 uses *source filtering*, which allows hosts to report interest in receiving multicast packets only from specific source addresses, or from all but specific source addresses.

The SI4091 supports the following IGMPv3 filter modes:

- INCLUDE mode: The host requests membership to a multicast group and provides a list of IPv4 addresses from which it wants to receive traffic.
- EXCLUDE mode: The host requests membership to a multicast group and provides a list of IPv4 addresses from which it *does not* want to receive traffic. This indicates that the host wants to receive traffic only from sources that are not part of the Exclude list. To disable snooping on EXCLUDE mode reports, use the following command:

```
SI 4091(config)# no ip igmp snoop igmpv3 exclude
```

By default, the SI4091 snoops the first eight sources listed in the IGMPv3 Group Record. Use the following command to change the number of snooping sources:

```
SI 4091(config)# ip igmp snoop igmpv3 sources <1-64>
```

IGMPv3 Snooping is compatible with IGMPv1 and IGMPv2 Snooping. You can disable snooping on version 1 and version 2 reports, using the following command:

```
SI 4091(config)# no ip igmp snoop igmpv3 v1v2
```

IGMP Snooping Configuration Example

This section provides steps to configure IGMP Snooping on the SI4091, using the Command-Line Interface (CLI).

1. Configure port and VLAN membership on the SI4091.
2. Add VLANs to IGMP Snooping and enable IGMP Snooping.

```
SI 4091(config)# ip igmp snoop vlan 1
SI 4091(config)# ip igmp snoop enable
```

3. Enable IGMPv3 Snooping (optional).

```
SI 4091(config)# ip igmp snoop igmpv3 enable
```

4. Enable IGMP.

```
SI 4091(config)# ip igmp enable          (Turn on IGMP)
```

5. View dynamic IGMP information.

   To display information about IGMP Groups:

```
SI 4091# show ip igmp groups

Total entries: 5  Total IGMP groups: 3
Note: The <Total IGMP groups> number is computed as
      the number of unique (Group, Vlan) entries!

Note: Local groups (224.0.0.x) are not snooped and will not appear.
     Source         Group       VLAN    Port   Version   Mode  Expires  Fwd
 -------------- --------------- ------- ------ --------  ----- -------  ---
 10.1.1.1       232.1.1.1        2      EXT4      V3      INC   4:16    Yes
 10.1.1.5       232.1.1.1        2      EXT4      V3      INC   4:16    Yes
    *           232.1.1.1        2      EXT4      V3      INC    -      No
 10.10.10.43    235.0.0.1        9      EXT1      V3      INC   2:26    Yes
    *           235.0.0.1        9      EXT1      V3      EXC    -      Yes
```

   To display information about Mrouters learned by the SI4091:

```
SI 4091# show ip igmp mrouter

Total entries: 3 Total number of dynamic mrouters: 2
Total number of installed static mrouters : 1

SrcIP              VLAN    Port       Version   Expires   MRT       QRV   QQIC

------------------ ------- -------    --------- --------  -------   ----  ----
10.1.1.1           2       EXT8          V3       4:09      128       2    125
10.1.1.5           3       EXT9          V2       4:09      125       -      -
10.10.10.43        9       EXT10         V2      static     -         -      -
```

Note: If IGMP Snooping v1/v2 is enabled and IGMPv3 Snooping is disabled, the output of IGMPv3 reports and queries show some items as IGMPv3 (`V3`), though they retain v2 behavior. For example, the `Source` IPv4 address is not relevant for v2 entries.

Static Multicast Router

A static multicast router (Mrouter) can be configured for a particular port on a particular VLAN.

A total of 128 static Mrouters can be configured on the SI4091. Both internal and external ports can accept a static Mrouter.

Note: When static Mrouters are used, the SI4091 will continue learning dynamic Mrouters via IGMP snooping. However, dynamic Mrouters may not replace static Mrouters. If a dynamic Mrouter has the same port and VLAN combination as a static Mrouter, the dynamic Mrouter will not be learned.

Following is an example of configuring a static multicast router:

1. For each Mrouter, configure a port, VLAN, and IGMP version of the multicast router.

```
SI 4091(config)# ip igmp mrouter EXT5 1 2
```

2. Verify the configuration.

```
SI 4091(config)# show ip igmp mrouter
```

FastLeave

In normal IGMP operation, when the SI4091 receives an IGMPv2 *leave* message, it sends a Group-Specific Query to determine if any other devices in the same group (and on the same port) are still interested in the specified multicast group traffic. The SI4091 removes the affiliated port from that particular group, if it does not receive an IGMP Membership Report within the query-response-interval.

With FastLeave enabled on the VLAN, a port can be removed immediately from the port list of the group entry when the IGMP Leave message is received, unless a multicast router was learned on the port.

Enable FastLeave only on VLANs that have only one host connected to each physical port.

IGMP Filtering

With IGMP Filtering, you can allow or deny a port to learn certain IGMP or IPMC groups. This allows you to restrict users from receiving certain multicast traffic.

If access to a multicast group is denied, IGMP Membership Reports from the port are dropped, and the port is not allowed to receive IPv4 multicast traffic from that group. If access to the multicast group is allowed, Membership Reports from the port are forwarded for normal processing.

To configure IGMP Filtering, you must globally enable IGMP filtering, define an IGMP filter, assign the filter to a port, and enable IGMP Filtering on the port. To define an IGMP filter, you must configure a range of IPv4 multicast groups, choose whether the filter will allow or deny multicast traffic for groups within the range, and enable the filter.

Configure the Range

Each IGMP Filter allows you to set a start and end point that defines the range of IPv4 addresses upon which the filter takes action. Each IPv4 address in the range must be between 224.0.0.0 and 239.255.255.255.

Configure the Action

Each IGMP filter can allow or deny IPv4 multicasts to the range of IPv4 addresses configured. If you configure the filter to deny IPv4 multicasts, then IGMP Membership Reports from multicast groups within the range are dropped. You can configure a secondary filter to allow IPv4 multicasts to a small range of addresses within a larger range that a primary filter is configured to deny. The two filters work together to allow IPv4 multicasts to a small subset of addresses within the larger range of addresses.

Note: Lower-numbered filters take precedence over higher-number filters. For example, the action defined for IGMP Filter 1 supersedes the action defined for IGMP Filter 2.

Configure IGMP Filtering

Follow these steps to configure IGMP filtering:

1. Enable IGMP filtering on the SI4091.

```
SI 4091(config)# ip igmp filtering
```

2. Define an IGMP filter with IPv4 information.

```
SI 4091(config)# ip igmp profile 1 range 224.0.0.0 226.0.0.0
SI 4091(config)# ip igmp profile 1 action deny
SI 4091(config)# ip igmp profile 1 enable
```

3. Assign the IGMP filter to a port.

```
SI 4091(config)# interface port EXT3
SI 4091(config-if)# ip igmp profile 1
SI 4091(config-if)# ip igmp filtering
```

Switch Behavior and IGMP Filtering

Without IGMP filtering configured and IPMC flooding enabled (the default configuration, all traffic is forwarded.

Without IGMP filtering configured and IPMC flooding disabled, no traffic is forwarded.

With a basic IGMP configuration with no multicast entries installed, whether or not IPMC flooding is enabled, traffic is forwarded to an IGMP mrouter port at the CoPP rate for destination-unknown packets and is not forwarded to a port where IGMP group entries are installed.

With a basic IGMP configuration with mrouters and groups, whether or not IPMC flooding is enabled, traffic is forwarded to an IGMP mrouter port at the CoPP rate for destination-unknown packets and is not forwarded to a port where IGMP group entries are installed.

Chapter 15. Service Location Protocol

Service Location Protocol (SLP) allows the switch to provide dynamic directory services that helps users find servers by attributes rather than by name or address. SLP eliminates the need for a user to know the name of a network host supporting a service. SLP allows the user to bind a service description to the network address of the service.

Service Location Protocol is described in RFC 2608.

Note: SLP is not supported on the internal management port (MGT).

SLP defines specialized components called agents that perform tasks and support services as follows:

- User Agent (UA) supports service query functions. It requests service information for user applications. The User Agent retrieves service information from the Service Agent or Directory Agents. A Host On-Demand client is an example of a User Agent.
- Service Agent (SA) provides service registration and service advertisement. **Note**: In this release, SA supports UA/DA on Linux with SLPv2 support.
- Directory Agent (DA) collects service information from Service Agents to provide a repository of service information in order to centralize it for efficient access by User Agents. There can only be one Directory Agent present per given host.

The Directory Agent acts as an intermediate tier in the SLP architecture, placed between the User Agents and the Service Agents, so they communicate only with the Directory Agent instead of with each other. This eliminates a large portion of the multicast request or reply traffic on the network, and it protects the Service Agents from being overwhelmed by too many service requests.

Services are described by the configuration of attributes associated with a type of service. A User Agent can select an appropriate service by specifying the attributes that it needs in a service request message. When service replies are returned, they contain a Uniform Resource Locator (URL) pointing to the service desired, and other information, such as server load, needed by the User Agent.

Active DA Discovery

When a Service Agent or User Agent initializes, it can perform Active Directory Agent Discovery using a multicast service request and specifies the special, reserved service type (`service:directory-agent`). Active DA Discovery is achieved through the same mechanism as any other discovery using SLP.

The Directory Agent replies with unicast service replies, which provides the URLs and attributes of the requested service.

SLP Configuration

Use the following CLI commands to configure SLP for the switch:

Table 16. *SLP ISCLI Commands*

Command Syntax and Usage
`[no] ip slp enable` Enables or disables SLP on the switch. **Command mode:** Global configuration
`[no] ip slp active-da-discovery enable` Enables or disables Active DA Discovery. **Command mode:** Global configuration
`ip slp active-da-discovery start-wait-time <1-10>` Configures the wait time before starting Active DA Discovery, in seconds. The default value is 3 seconds. **Command mode:** Global configuration
`clear ip slp directory-agents` Clears all Directory Agents learned by the switch. **Command mode:** Global configuration
`show ip slp information` Displays SLP information. **Command mode:** All
`show ip slp directory-agents` Displays Directory Agents learned by the switch. **Command mode:** All
`show ip slp user-agents` Displays User Agents information. **Command mode:** All
`show ip slp counters` Displays SLP statistics. **Command mode:** All
`clear ip slp counters` Clears all Directory Agents learned by the switch. **Command mode:** Global configuration

Part 5: Extended Features

Chapter 16. CEE and FCoE

This chapter provides conceptual background and configuration examples for using Converged Enhanced Ethernet (CEE) features of the SI4091 System Interconnect Module, with an emphasis on Fibre Channel over Ethernet (FCoE) solutions. The following topics are addressed in this chapter:

- "Fibre Channel over Ethernet" on page 198

  Fibre Channel over Ethernet (FCoE) allows Fibre Channel traffic to be transported over Ethernet links. This provides an evolutionary approach toward network consolidation, allowing Fibre Channel equipment and tools to be retained, while leveraging cheap, ubiquitous Ethernet networks for growth.

  - "FCoE Initialization Protocol Snooping" on page 204

    Using FCoE Initialization Protocol (FIP) snooping, the SI4091 examines the FIP frames exchanged between ENodes and FCFs. This information is used to dynamically determine the ACLs required to block certain types of undesired or unvalidated traffic on FCoE links.

- "Converged Enhanced Ethernet" on page 201

  Converged Enhanced Ethernet (CEE) refers to a set of IEEE standards developed primarily to enable FCoE, requiring enhancing the existing Ethernet standards to make them lossless on a per-priority traffic basis, and providing a mechanism to carry converged (LAN/SAN/IPC) traffic on a single physical link. CEE features can also be utilized in traditional LAN (non-FCoE) networks to provide lossless guarantees on a per-priority basis, and to provide efficient bandwidth allocation.

  - "Priority-Based Flow Control" on page 209

    Priority-Based Flow Control (PFC) extends 802.3x standard flow control to allow the switch to pause traffic based on the 802.1p priority value in each packet's VLAN tag. PFC is vital for FCoE environments, where SAN traffic must remain lossless and must be paused during congestion, while LAN traffic on the same links is delivered with "best effort" characteristics.

  - "Enhanced Transmission Selection" on page 213

    Enhanced Transmission Selection (ETS) provides a method for allocating link bandwidth based on the 802.1p priority value in each packet's VLAN tag. Using ETS, different types of traffic (such as LAN, SAN, and management) that are sensitive to different handling criteria can be configured either for specific bandwidth characteristics, low-latency, or best-effort transmission, despite sharing converged links as in an FCoE environment.

  - "Data Center Bridging Capability Exchange" on page 219

    Data Center Bridging Capability Exchange Protocol (DCBX) allows neighboring network devices to exchange information about their capabilities. This is used between CEE-capable devices for the purpose of discovering their peers, negotiating peer configurations, and detecting misconfigurations.

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) is an effort to converge two of the different physical networks in today's data centers. It allows Fibre Channel traffic (such as that commonly used in Storage Area Networks, or SANs) to be transported without loss over 10Gb Ethernet links (typically used for high-speed Local Area Networks, or LANs). This provides an evolutionary approach toward network consolidation, allowing Fibre Channel equipment and tools to be retained, while leveraging cheap, ubiquitous Ethernet networks for growth.

With server virtualization, servers capable of hosting both Fibre Channel and Ethernet applications will provide advantages in server efficiency, particularly as FCoE-enabled network adapters provide consolidated SAN and LAN traffic capabilities.

The SI4091 with ENOS 8.4 software is compliant with the INCITS T11.3, FC-BB-5 FCoE specification and supports up to 2048 FCoE login sessions.

The FCoE Topology

In an end-to-end Fibre Channel network, switches and end devices generally establish trusted, point-to-point links. Fibre Channel switches validate end devices, enforce zoning configurations and device addressing, and prevent certain types of errors and attacks on the network.

In a converged multi-hop FCoE network where Fibre Channel devices are bridged to Ethernet devices, the direct point-to-point assurances normally provided by the Fibre Channel fabric may be lost in the transition between the different network types. However, the SI4091 provides a solution for this problem.

Figure 18. A Mixed Fibre Channel and FCoE Network

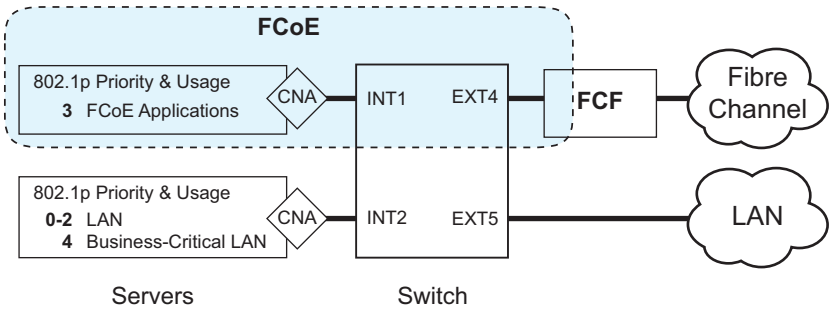

In Figure 18, the Fibre Channel network is connected to the FCoE network through a Fibre Channel Forwarder (FCF) bridge. The FCF acts as a Fibre Channel gateway to and from the multi-hop FCoE network.

For the FCoE portion of the network, the FCF is connected to the FCoE-enabled SI4091, which is internally connected to a blade server (running Fibre Channel applications) through an FCoE-enabled Converged Network Adapter (CNA) known in Fibre Channel as an Ethernet Node (ENode).

Note: The figure also shows a non-FCoE LAN server connected to the SI4091 using a CNA. This allows the LAN server to take advantage of some CEE features that are useful even outside of an FCoE environment.

FCoE Security

The SI4091 can block undesired or unvalidated traffic on FCoE links that exists outside the regular Fibre Channel topology. The Ethernet ports used in FCoE are configured with Access Control Lists (ACLs) that are narrowly tailored to permit expected FCoE traffic to and from confirmed FCFs and ENodes, and deny all other FCoE or FCoE Initialization Protocol (FIP) traffic. This ensures that all FCoE traffic to and from the ENode passes through the FCF.

Because manual ACL configuration is an administratively complex task, the SI4091 can automatically and dynamically configure the ACLs required for use with FCoE. Using FIP snooping (see "FCoE Initialization Protocol Snooping" on page 204), the SI4091 examines the FIP frames normally exchanged between the FCF and ENodes to determine information about connected FCoE devices. This information is used to automatically determine the appropriate ACLs required to block certain types of undesired or unvalidated FCoE traffic.

Automatic FCoE-related ACLs are independent from ACLs used for typical Ethernet purposes.

The FIP snooping feature is turned off by default. When the feature is turned on, FIP snooping is enabled on all ports by default, but can be individually enabled or disabled on port by port basis.

FCoE Requirements

The following are required for implementing FCoE using the SI4091 with ENOS 8.4 software:

- The SI4091 must be connected to the Fibre Channel network through an FCF such as a Lenovo RackSwitch G8264CS or a Cisco Nexus 5000 Series Switch.
- For each SI4091 internal port participating in FCoE, the connected blade server must use the supported FCoE CNA. Emulex Virtual Fabric Adapter 2-port 10Gb LOM and Emulex Virtual Fabric Adapter (Fabric Mezz) for Lenovo Flex System, which includes vNIC support (with some additional topology rules), is currently supported.
- For each SI4091 internal port participating in FCoE, the connected blade server must include the appropriate FCoE licenses installed, as obtained using the IBM website Features on Demand (FoD) service. Contact your sales representative for more information on obtaining server feature licenses.
- CEE must be turned on (see "Turning CEE On or Off" on page 201). When CEE is on, the DCBX, PFC, and ETS features are enabled and configured with default FCoE settings. These features may be reconfigured, but must remain enabled in order for FCoE to function.
- FIP snooping must be turned on (see "FCoE Initialization Protocol Snooping" on page 204). When FIP snooping is turned on, the feature is enabled on all ports by default. The administrator can disable FIP snooping on individual ports that do not require FCoE, but FIP snooping must remain enabled on all FCoE ports to support secured FCoE sessions.
- An FCOE VLAN must always be set as tagged and never set as a native VLAN.

FCoE and Aggregation

Enterprise NOS 8.4 supports aggregation for FCoE connections. The Link Aggregation Group (LAG) can be used for separate FCoE traffic, or for Ethernet and FCoE traffic. Ports directly connected to servers cannot be combined in a LAG group.

Uplink ports, connected to the FCF, can be grouped as static or dynamic LAGs.

Internal ports cannot be grouped as LAGs.

Normal LAG operations such as creating/enabling the LAG, and adding/removing member ports can be performed. When a port is added to a LAG, FCFs previously detected on the port will be deleted. The deleted FCF may be relearned later. However, this may cause flickering in the network traffic.

Priority-based Flow Control (PFC), and Data Center Bridging (DCBX) are configured on a per-port basis. Each port in a LAG must have the same PFC, and DCBX configuration. When a port ceases to be the LAG member, its configuration does not change.

Converged Enhanced Ethernet

Converged Enhanced Ethernet (CEE) refers to a set of IEEE standards designed to allow different physical networks with different data handling requirements to be converged together, simplifying management, increasing efficiency and utilization, and leveraging legacy investments without sacrificing evolutionary growth.

CEE standards were developed primarily to enable Fibre Channel traffic to be carried over Ethernet networks. This required enhancing the existing Ethernet standards to make them lossless on a per-priority traffic basis, and to provide a mechanism to carry converged (LAN/SAN/IPC) traffic on a single physical link. Although CEE standards were designed with FCoE in mind, they are not limited to FCoE installations. CEE features can be utilized in traditional LAN (non-FCoE) networks to provide lossless guarantees on a per-priority basis, and to provide efficient bandwidth allocation based on application needs.

Turning CEE On or Off

By default on the SI4091, CEE is turned off. To turn CEE on or off, use the following command:

```
SI 4091(config)# [no] cee enable
```

For an example, see "FIP Snooping Configuration" on page 208.

Note:

- Turning CEE on will automatically change some 802.1p QoS and 802.3x standard flow control settings on the SI4091. Read the following material carefully to determine whether you will need to take action to reconfigure expected settings.
- It is recommended that you backup your configuration prior to turning CEE on. Viewing the file will allow you to manually re-create the equivalent configuration once CEE is turned on, and will also allow you to recover your prior configuration if you need to turn CEE off.

Effects on Link Layer Discovery Protocol

When CEE is turned on, Link Layer Discovery Protocol (LLDP) is automatically turned on and enabled for receiving and transmitting DCBX information. LLDP cannot be turned off while CEE is turned on.

Effects on 802.1p Quality of Service

When CEE is off (the default), the SI4091 allows 802.1p priority values to be used for Quality of Service (QoS) configuration. The 802.1p QoS default settings are shown in Table 17 on page 202, but can be changed by the administrator.

While CEE is turned on, 802.1p QoS is replaced by ETS (see "Enhanced Transmission Selection" on page 213). As a result, while CEE is turned on, the 802.1p QoS configuration commands are no longer available on the SI4091. The commands are restored when CEE is turned off.

In addition, while CEE is turned on, any prior 802.1p QoS settings are replaced with new defaults designed for use with ETS priority groups (PGIDs) as shown in Table 17:

Table 17. *CEE Effects on 802.1p Defaults*

802.1p QoS Configuration With CEE Off			ETS Configuration With CEE On		
Priority	**COSq**	**Weight**	**Priority**	**COSq**	**PGID**
0	0	1	0	0	0
1	1	2	1	0	0
2	2	3	2	0	0
3	3	4	3	1	1
4	4	5	4	2	2
5	5	7	5	2	2
6	6	15	6	2	2
7	7	0	7	2	2

When CEE is on, the default ETS configuration also allocates a portion of link bandwidth to each PGID as shown in Table 18:

Table 18. *Default ETS Bandwidth Allocation*

PGID	Typical Use	Bandwidth
0	LAN	10%
1	SAN	50%
2	Latency-sensitive LAN	40%

If the prior, non-CEE configuration used 802.1p priority values for different purposes, or does not expect bandwidth allocation as shown in Table 18 on page 202, when CEE is turned on, the administrator should reconfigure ETS settings as appropriate.

Effects on Flow Control

When CEE is off (the default), 802.3x standard flow control is enabled on all SI4091 ports by default.

When CEE is turned on, standard flow control is disabled on all ports, and in its place, PFC (see "Priority-Based Flow Control" on page 209) is enabled on all ports for 802.1p priority value 3. This default is chosen because priority value 3 is commonly used to identify FCoE traffic in a CEE environment and must be guaranteed lossless behavior. PFC is disabled for all other priority values.

Until rebooted, each time CEE is turned off, the prior 802.3x standard flow control settings will be restored (including any previous changes from the defaults). Conversely, each time CEE is turned on, the previously configured PFC settings are restored. Once the switch is rebooted with CEE turned on, prior 802.3x standard flow control settings are returned to factory default values. Changes to standard flow control settings will be lost.

It is recommend that a configuration backup be made prior to turning CEE on or off. Viewing the configuration file will allow the administrator to manually re-create the equivalent configuration under the new CEE mode, and will also allow for the recovery of the prior configuration if necessary.

When CEE is on, PFC can be enabled only on priority value 3 and one other priority. If flow control is required on additional priorities on any given port, consider using standard flow control on that port, so that regardless of which priority traffic becomes congested, a flow control frame is generated.

FCoE Initialization Protocol Snooping

FCoE Initialization Protocol (FIP) snooping is an FCoE feature. In order to enforce point-to-point links for FCoE traffic outside the regular Fibre Channel topology, Ethernet ports used in FCoE can be automatically and dynamically configured with Access Control Lists (ACLs).

Using FIP snooping, the SI4091 examines the FIP frames normally exchanged between the FCF and ENodes to determine information about connected FCoE devices. This information is used to create narrowly tailored ACLs that permit expected FCoE traffic to and from confirmed Fibre Channel nodes, and deny all other undesirable FCoE or FIP traffic.

In case of LAGs, FIP traffic from a particular FCF can be received by any member port on which the FCF was detected.

Global FIP Snooping Settings

By default, the FIP snooping feature is turned off for the SI4091. The following commands are used to turn the feature on or off:

```
SI 4091(config)# [no] fcoe fips enable
```

Note: FIP snooping requires CEE to be turned on (see "Turning CEE On or Off" on page 201).

When FIP snooping is on, port participation may be configured on a port-by-port basis (see "FIP Snooping for Specific Ports" on page 205).

Before turning on the FIP snooping feature, you must manually disable FIP snooping for internal and external ports where FCoE traffic is not desired.

When FIP snooping is off, all FCoE-related ACLs generated by the feature are removed from all SI4091 ports.

FIP snooping configuration must be the same on all member ports in a LAG. If the configuration of a member port is changed, an error message, similar to the following, will be displayed.

```
FAIL: Trunk <ID>  FIP Snooping port <number> and port <number> need to have
the same fips config.
```

The FIP snooping configuration changes must be manually applied to all member ports.

For an example, see "FIP Snooping Configuration" on page 208.

FIP Snooping for Specific Ports

When FIP snooping is globally turned on (see above), ports may be individually configured for participation in FIP snooping and automatic ACL generation. By default, FIP snooping is enabled for each port. To change the setting for any specific port, use the following CLI commands:

```
SI 4091(config)# [no] fcoe fips port <port alias or number> enable
```

When FIP snooping is enabled on a port, FCoE-related ACLs will be automatically configured.

When FIP snooping is disabled on a port, all FCoE-related ACLs on the port are removed, and the SI4091 will enforce no FCoE-related rules for traffic on the port.

Before turning on the FIP snooping feature, you must manually disabled FIP snooping for internal and external ports where FCoE traffic is not desired.

Note: FIP Snooping and IPv6 ACLs are not supported simultaneously on the same ports. To use FIP snooping, remove IPv6 ACLs from the port.

Port FCF and ENode Detection

When FIP snooping is enabled on a port, the port is placed in FCF auto-detect mode by default. In this mode, the port assumes connection to an ENode unless FIP packets show the port is connected to an FCF.

Ports can also be specifically configured as to whether automatic FCF detection should be used, or whether the port is connected to an FCF or ENode:

```
SI 4091(config)# fcoe fips port <port alias or number> fcf-mode {auto|on|off}
```

When FCF mode is `on`, the port is assumed to be connected to a trusted FCF, and only ACLs appropriate to FCFs will be installed on the port. When `off`, the port is assumed to be connected to an ENode, and only ACLs appropriate to ENodes will be installed. When the mode is changed (either through manual configuration or as a result of automatic detection), the appropriate ACLs are automatically added, removed, or changed to reflect the new FCF or ENode connection.

FCoE Connection Timeout

FCoE-related ACLs are added, changed, and removed as FCoE device connection and disconnection are discovered. In addition, the administrator can enable or disable automatic removal of ACLs for FCFs and other FCoE connections that timeout (fail or are disconnected) without FIP notification.

By default, automatic removal of ACLs upon timeout is enabled. To change this function, use the following CLI commands:

```
SI 4091(config)# [no] fcoe fips timeout-acl
```

FCoE ACL Rules

When FIP Snooping is enabled on a port, the SI4091 automatically installs the appropriate ACLs to enforce the following rules for FCoE traffic:

- Ensure that FIP frames from ENodes may only be addressed to FCFs.
- Flag important FIP packets for processing.
- Ensure no end device uses an FCF MAC address as its source.
- Each FCoE port is assumed to be connected to an ENode and includes ENode-specific ACLs installed, until the port is either detected or configured to be connected to an FCF.
- Ports that are configured to have FIP snooping disabled will not have any FIP or FCoE related ACLs installed.
- Prevent transmission of all FCoE frames from an ENode prior to its successful completion of login (FLOGI) to the FCF.
- After successful completion of FLOGI, ensure that the ENode uses only those FCoE source addresses assigned to it by FCF.
- After successful completion of FLOGI, ensure that all ENode FCoE source addresses originate from or are destined to the appropriate ENode port.
- After successful completion of each FLOGI, ensure that FCoE frames may only be addressed to the FCFs that accept them.

Initially, a basic set of FCoE-related ACLs will be installed on all ports where FIP snooping is enabled. As the SI4091 encounters FIP frames and learns about FCFs and ENodes that are attached or disconnect, ACLs are dynamically installed or expanded to provide appropriate security.

When an FCoE connection logs out, or times out (if ACL timeout is enabled), the related ACLs will be automatically removed.

FCoE-related ACLs are independent of manually configured ACLs used for regular Ethernet purposes. FCoE ACLs generally have a higher priority over standard ACLs.

FCoE VLANs

Before the SI4091 applies FIP Snooping, all internal SI4091 ports connected to ENodes and all external ports connected to FCFs should be members of at least one common VLAN (for example, default VLAN 1). This allows the ENode CNA and the FCF to exchange initial FIP VLAN request and notification packets. Once FIP Snooping is applied, FCoE packets are exchanged using one configured FCoE VLAN for each attached FCF.

Each ENode port must have VLAN tagging enabled, and must belong to the same VLAN as the FCF to which it will connect. In topologies where a single FCF is connected to the SI4091, all ENode and FCF ports belong to the same VLAN (typically VLAN 1002). When multiple FCFs are connected to the SI4091, each FCF must be assigned a unique VLAN, and each ENode must be assigned to the VLAN for only one particular FCF.

The administrator must ensure that the VLAN configured for each FCF and its ENodes is supported by the participating FCF and ENode CNAs.

Note: FCoE ports can have any native VLAN. A Fibre Channel VLAN must not be configured as the native VLAN for FCoE ports.

Viewing FIP Snooping Information

ACLs automatically generated under FIP snooping are independent of regular, manually configure ACLs, and are not listed with regular ACLs in SI4091 information and statistics output. Instead, FCoE ACLs are shown using the following CLI commands:

```
SI 4091# show fcoe fips information    (Show all FIP-related information)

   -or-

SI 4091# show fcoe fips port <port alias or number> information(Show FIP info for a
selected port)
```

For example:

```
SI 4091# show fcoe fips port ext4 information

FIP Snooping on port ext4:
This port has been detected to be an FCF port.

FIPS ACLs configured on this port:
Ethertype 0x8914, action permit.
dmac 00:00:18:01:00:XX, Ethertype 0x8914, action permit.
```

For each ACL, the required traffic criteria are listed, along with the action taken (permit or deny) for matching traffic. ACLs are listed in order of precedence and evaluated in the order shown.

The administrator can also view other FCoE information:

```
SI 4091# show fcoe fips fcf            (Show all detected FCFs)

   -or-

SI 4091# show fcoe fips fcoe           (Show all FCoE connections)
```

Manually Removing FIP Entries

Use the following operational command to delete a specific FCF entry and all associated ACLs from the switch:

```
SI 4091# no fcoe fips fcf <FCF MAC address> [<VLAN number>]
```

FIP Snooping Configuration

In this example (as shown in Figure 18 on page 198), FCoE devices are connected to port EXT4 for the FCF bridge connection, and INT1 for an ENode. FIP snooping can be configured on these ports using the following CLI commands:

1. Turn CEE on:

```
SI 4091(config)# cee enable
```

Note: Turning CEE on or off will automatically change some 802.1p QoS and 802.3x standard flow control settings (see "Turning CEE On or Off" on page 201).

2. Disable FIP snooping on all non-FCoE external ports:

```
SI 4091(config)# no fcoe fips port ext1-ext3,ext5-ext10 enable
```

Note: FIP snooping must remain enabled for all FCoE ports.

3. Turn global FIP snooping on:

```
SI 4091(config)# fcoe fips enable
```

4. Configure the FCoE VLAN:

```
SI 4091(config)# vlan 1002                    (Select and enable the VLAN)
SI 4091(config-vlan)# exit                    (Exit VLAN configuration mode)
```

5. Configure the FCoE ports and enable VLAN tagging.

   Place all FCoE ports associated with each FCF into a unique VLAN supported by that FCF:

```
SI 4091(config)# interface port INT1,EXT4          (Select FCoE ports)
SI 4091(config-if)# switchport mode trunk          (Enable VLAN tagging)
SI 4091(config-if)# switchport trunk allowed vlan 1,1002
SI 4091(config-if)# exit                           (Exit port configuration mode)
```

Note: Placing ports into the VLAN *after* tagging is enabled helps to ensure that their port VLAN ID (PVID) is not accidentally changed.

6. Enable FIP snooping on FCoE ports (set by default), and set the desired FCF mode:

```
SI 4091(config)# fcoe fips port INT1,EXT4 enable (Enable FIPS on FCoE ports)
SI 4091(config)# fcoe fips port INT1 fcf-mode off(Set as ENode connection)
SI 4091(config)# fcoe fips port EXT4 fcf-mode on (Set as FCF connection)
```

7. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

Priority-Based Flow Control

Priority-based Flow Control (PFC) is defined in IEEE 802.1Qbb. PFC extends the IEEE 802.3x standard flow control mechanism. Under standard flow control, when a port becomes busy, the SI4091 manages congestion by pausing all the traffic on the port, regardless of the traffic type. PFC provides more granular flow control, allowing the SI4091 to pause specified types of traffic on the port, while other traffic on the port continues.

PFC pauses traffic based on 802.1p priority values in the VLAN tag. The administrator can assign different priority values to different types of traffic and then enable PFC for up to two specific priority values: priority value 3, and one other. The configuration can be applied on a port-by-port basis, or globally for all ports on the SI4091. Then, when traffic congestion occurs on a port (caused when ingress traffic exceeds internal buffer thresholds), only traffic with priority values where PFC is enabled is paused. Traffic with priority values where PFC is disabled proceeds without interruption but may be subject to loss if port ingress buffers become full.

Although PFC is useful for a variety of applications, it is required for FCoE implementation where storage (SAN) and networking (LAN) traffic are converged on the same Ethernet links. Typical LAN traffic tolerates Ethernet packet loss that can occur from congestion or other factors, but SAN traffic must be lossless and requires flow control.

For FCoE, standard flow control would pause both SAN and LAN traffic during congestion. While this approach would limit SAN traffic loss, it could degrade the performance of some LAN applications that expect to handle congestion by dropping traffic. PFC resolves these FCoE flow control issues. Different types of SAN and LAN traffic can be assigned different IEEE 802.1p priority values. PFC can then be enabled for priority values that represent SAN and LAN traffic that must be paused during congestion, and disabled for priority values that represent LAN traffic that is more loss-tolerant.

PFC requires CEE to be turned on ("Turning CEE On or Off" on page 201). When CEE is turned on, PFC is enabled on priority value 3 by default. Optionally, the administrator can also enable PFC on one other priority value, providing lossless handling for another traffic type, such as for a business-critical LAN application.

Note: For any given port, only one flow control method can be implemented at any given time: either PFC or standard IEEE 802.3x flow control.

Global vs. Port-by-Port PFC Configuration

PFC requires CEE to be turned on (see "Turning CEE On or Off" on page 201). When CEE is turned on, standard flow control is disabled on all ports, and PFC is enabled on all ports for 802.1p priority value 3. This default is chosen because priority value 3 is commonly used to identify FCoE traffic in a CEE environment and must be guaranteed lossless behavior. PFC is disabled for all other priority values by default, but can be enabled for one additional priority value.

The administrator can also configure PFC on a port-by-port basis. The method used will typically depend on the following:

- Port-by-port PFC configuration is desirable in most mixed environments where some SI4091 ports are connected to CEE-capable (FCoE) switches, gateways, and Converged Network Adapters (CNAs), and other SI4091 ports are connected to non-CEE Layer 2/Layer 3 switches, routers and Network Interface Cards (NICs).
- Global PFC configuration is preferable in networks that implement end-to-end CEE devices. For example, if all ports are involved with FCoE and can use the same SAN and LAN priority value configuration with the same PFC settings, global configuration is easy and efficient.
- Global PFC configuration can also be used in some mixed environments where traffic with PFC-enabled priority values occurs only on ports connected to CEE devices, and not on any ports connected to non-CEE devices. In such cases, PFC can be configured globally on specific priority values even though not all ports make use them.
- PFC is not restricted to CEE and FCoE networks. In any LAN where traffic is separated into different priorities, PFC can be enabled on priority values for loss-sensitive traffic. If all ports have the same priority definitions and utilize the same PFC strategy, PFC can be globally configured.

Note: When using global PFC configuration in conjunction with the ETS feature (see "Enhanced Transmission Selection" on page 213), ensure that only pause-tolerant traffic (such as lossless FCoE traffic) is assigned priority values where PFC is enabled. Pausing other types of traffic can have adverse effects on LAN applications that expect uninterrupted traffic flow and tolerate dropping packets during congestion. Use PFC globally only if all priority values assigned for lossless traffic on one or more ports does not carry loss-tolerant traffic on other ports.

PFC Configuration Example

Note: DCBX may be configured to permit sharing or learning PFC configuration with or from external devices. This example assumes that PFC configuration is being performed manually. See "Data Center Bridging Capability Exchange" on page 219 for more information on DCBX.

This example is consistent with the network shown in Figure on page 198. In this example, the following topology is used.

Table 19. *Port-Based PFC Configuration*

SI4091 Port	802.1p Priority	Usage	PFC Disabled	PFC Enabled
EXT5	0-2	LAN	Disabled	
	3	*(not used)*		Enabled
	4	Business-critical LAN		Enabled
	others	*(not used)*	Disabled	
EXT4	3	FCoE (to FCF bridge)		Enabled
	others	*(not used)*	Disabled	
INT1	3	FCoE		Enabled
	others	*(not used)*	Disabled	
INT2	0-2	LAN	Disabled	
	3	*(not used)*		Enabled
	4	Business-critical LAN		Enabled
	others	*(not used)*	Disabled	

In this example, PFC is to facilitate lossless traffic handling for FCoE (priority value 3) and a business-critical LAN application (priority value 4).

The example topology shown in Table 19 on page 211 can be configured using the following commands:

1. Turn CEE on.

```
SI 4091(config)# cee enable
```

Note: Turning CEE on or off will automatically change some 802.1p QoS and 802.3x standard flow control settings (see "Turning CEE On or Off" on page 201).

2. Enable PFC for the FCoE traffic, if necessary.

Note: By default, PFC is enabled on priority 3 for all ports. If using the default settings, the commands in this step are not necessary and are shown here merely as an example of manual configuration.

```
SI 4091(config)# cee port INT1 pfc     (Turn on PFC for port INT1)
SI 4091(config)# cee port INT1 pfc priority 3 enable
                                       (Select and enable FCoE priority 3 PFC for the port)
SI 4091(config)# cee port INT1 pfc priority 3 description "FCoE"
                                       (Set priority description—optional)

SI 4091(config)# cee port EXT4 pfc     (Select and enable FCoE for port EXT4)
SI 4091(config)# cee port EXT4 pfc priority 3 enable
                                       (Select and enable FCoE priority 3 PFC for the port)
SI 4091(config)# cee port EXT4 pfc priority 3 description "FCoE"
                                       (Set priority description—optional)
```

3. Enable PFC for the business-critical LAN application:

```
SI 4091(config)# cee port INT2 pfc     (Turn on PFC for port INT2)
SI 4091(config)# cee port INT2 pfc priority 4 enable
                                       (Select and enable FCoE priority 4 PFC for the port)
SI 4091(config)# cee port INT2 pfc priority 4 description
"Business-critical LAN"
                                       (Set priority description—optional)

SI 4091(config)# cee port EXT5 pfc     (Select and enable FCoE for port EXT5)
SI 4091(config)# cee port EXT5 pfc priority 4 enable
                                       (Select and enable FCoE priority 4 PFC for the port)
SI 4091(config)# cee port EXT5 pfc priority 4 description
"Business-critical LAN"
                                       (Set priority description—optional)
```

Notes:

- By default, PFC is enabled for all ports. If the default settings are in place, the two `cee port` <*number*> `pcf` commands may be omitted. The commands are shown here as a manual configuration example only.
- Although PFC is enabled on the ports by default, it is not enabled on priority 4 by default. The `cee port` <*number*> `pfc priority` <*PGID*> `enable` commands are still required.

4. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) is defined in IEEE 802.1Qaz. ETS provides a method for allocating port bandwidth based on 802.1p priority values in the VLAN tag. Using ETS, different amounts of link bandwidth can specified for different traffic types (such as for LAN, SAN, and management).

ETS is an essential component in a CEE environment that carries different types of traffic, each of which is sensitive to different handling criteria, such as Storage Area Networks (SANs) that are sensitive to packet loss, and LAN applications that may be latency-sensitive. In a single converged link, such as when implementing FCoE, ETS allows SAN and LAN traffic to coexist without imposing contrary handling requirements upon each other.

The ETS feature requires CEE to be turned on (see "Turning CEE On or Off" on page 201).

802.1p Priority Values

Under the 802.1p standard, there are eight available priority values, with values numbered 0 through 7, which can be placed in the priority field of the 802.1Q VLAN tag:

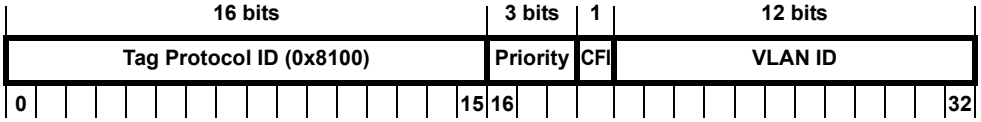

Servers and other network devices may be configured to assign different priority values to packets belonging to different traffic types (such as SAN and LAN).

ETS uses the assigned 802.1p priority values to identify different traffic types. The various priority values are assigned to priority groups (PGID), and each priority group is assigned a portion of available link bandwidth.

Priorities values within in any specific ETS priority group are expected to have similar traffic handling requirements with respect to latency and loss.

802.1p priority values may be assigned by the administrator for a variety of purposes. However, when CEE is turned on, the SI4091 sets the initial default values for ETS configuration as follows:

Figure 19. Default ETS Priority Groups

Typical Traffic Type	802.1p Priority	PGID	Bandwidth Allocation
LAN	0	0	10%
LAN	1		
LAN	2		
SAN	3	1	50%
Latency-Sensitive LAN	4	2	40%
Latency-Sensitive LAN	5		
Latency-Sensitive LAN	6		
Latency-Sensitive LAN	7		

In the assignment model shown in Figure 19 on page 213, priorities values 0 through 2 are assigned for regular Ethernet traffic, which has “best effort” transport characteristics.

Priority 3 is typically used to identify FCoE (SAN) traffic.

Priorities 4-7 are typically used for latency sensitive traffic and other important business applications. For example, priority 4 and 5 are often used for video and voice applications such as IPTV, Video on Demand (VoD), and Voice over IP (VoIP). Priority 6 and 7 are often used for traffic characterized with a “must get there” requirement, with priority 7 used for network control which is requires guaranteed delivery to support configuration and maintenance of the network infrastructure.

Note: The default assignment of 802.1p priority values on the SI4091 changes depending on whether CEE is on or off. See “Turning CEE On or Off” on page 201 for details.

Priority Groups

For ETS use, each 801.2p priority value is assigned to a priority group which can then be allocated a specific portion of available link bandwidth. To configure a priority group, the following is required:

- CEE must be turned on (see “Turning CEE On or Off” on page 201) for the ETS feature to function.
- Each priority group must be assigned a priority group ID (PGID), one or more 802.1p priority values, and allocated link bandwidth greater than 0%.

PGID

Each priority group is identified with number (0 through 7, and 15) known as the PGID.

PGID 0 through 7 may each be assigned a portion of the available bandwidth.

PGID 8 through 14 are reserved as per the 802.1Qaz ETS standard.

PGID 15 is a strict priority group. It is generally used for critical traffic, such as network management. Any traffic with priority values assigned to PGID 15 is permitted as much bandwidth as required, up to the maximum available on the SI4091. After serving PGID 15, any remaining link bandwidth is shared among the other groups, divided according to the configured bandwidth allocation settings.

All 802.1p priority values assigned to a particular PGID should have similar traffic handling requirements. For example, PFC-enabled traffic should not be grouped with non-PFC traffic. Also, traffic of the same general type should be assigned to the same PGID. Splitting one type of traffic into multiple 802.1p priorities, and then assigning those priorities to different PGIDs may result in unexpected network behavior.

Each 802.1p priority value may be assigned to only one PGID. However, each PGID may include multiple priority values. Up to eight PGIDs may be configured at any given time.

Assigning Priority Values to a Priority Group

Each priority group may be configured from its corresponding ETS Priority Group and assigned 802.1p priority values using the following command:

```
SI 4091(config)# cee global ets priority-group pgid <priority group number (0-7, or 15)> prio <priority list>
```

where *priority list* is one or more 802.1p priority values (with each separated by a space). For example, to assign priority values 0 through 2 to PGID 0:

```
SI 4091(config)# cee global ets priority-group pgid 0 prio 0 1 2
```

Note: Within any specific PGID, the PFC settings (see "Priority-Based Flow Control" on page 209) should be the same (enabled or disabled) for all priority values within the group. PFC can be enabled only on priority value 3 and one other priority. If the PFC setting is inconsistent within a PGID, a warning message is reported when attempting to apply the configuration.

When assigning priority values to a PGID, the specified priority value will be automatically removed from its old group and assigned to the new group when the configuration is applied.

Each priority value must be assigned to a PGID. Priority values may not be deleted or unassigned. To remove a priority value from a PGID, it must be moved to another PGID.

For PGIDs 0 through 7, bandwidth allocation can be configured through the ETS Priority Group. See for "Allocating Bandwidth" on page 216 for details.

Deleting a Priority Group

A priority group is automatically deleted when it contains no associated priority values, and its bandwidth allocation is set to 0%.

Note: The total bandwidth allocated to PGID 0 through 7 must equal exactly 100%. Reducing the bandwidth allocation of any group will require increasing the allocation to one or more of the other groups (see "Allocating Bandwidth" on page 216).

Allocating Bandwidth

Allocated Bandwidth for PGID 0 Through 7

The administrator may allocate a portion of the available switch bandwidth to PGIDs 0 through 7. Available bandwidth is defined as the amount of link bandwidth that remains after priorities within PGID 15 are serviced (see "Unlimited Bandwidth for PGID 15" on page 217), and assuming that all PGIDs are fully subscribed. If any PGID does not fully consume its allocated bandwidth, the unused portion is made available to the other priority groups.

Priority group bandwidth allocation can be configured using the following command:

```
SI 4091(config)# cee global ets priority-group [pgid 1 bandwidth <allocation
   >] [pgid 2 bandwidth <allocation >] [pgid 3 bandwidth <allocation >] [pgid 4
   bandwidth <allocation >] [pgid 5 bandwidth <allocation >] [pgid 6 bandwidth
   <allocation >] [pgid 7 bandwidth <allocation >] [pgid 8 bandwidth <allocation >]
```

where *allocation* represents the percentage of link bandwidth, specified as a number between 0 and 100, in 1% increments.

Note: You may omit the optional command segments for any priority groups where the allocation value is unchanged.

The following bandwidth allocation rules apply:

- Bandwidth allocation must be 0% for any PGID that has no assigned 802.1p priority values.
- Any PGID assigned one or more priority values must have a bandwidth allocation greater than 0%.
- Total bandwidth allocation for groups 0 through 7 must equal exactly 100%. Increasing or reducing the bandwidth allocation of any PGID also requires adjusting the allocation of other PGIDs to compensate.

If these conditions are not met, the SI4091 will report an error when applying the configuration.

To achieve a balanced bandwidth allocation among the various priority groups, packets are scheduled according to a weighted deficit round-robin (WDRR) algorithm. WDRR is aware of packet sizes, which can vary significantly in a CEE environment, making WDRR more suitable than a regular weighted round-robin (WRR) method, which selects groups based only on packet counts.

Note: Actual bandwidth used by any specific PGID may vary from configured values by up to 10% of the available bandwidth in accordance with 802.1Qaz ETS standard. For example, a setting of 10% may be served anywhere from 0% to 20% of the available bandwidth at any given time.

Unlimited Bandwidth for PGID 15

PGID 15 is permitted unlimited bandwidth and is generally intended for critical traffic (such as SI4091 management). Traffic in this group is given highest priority and is served before the traffic in any other priority group.

If PGID 15 has low traffic levels, most of the system's bandwidth will be available to serve priority groups 0 through 7. However, if PGID 15 consumes a larger part of the system's total bandwidth, the amount available to the other groups is reduced.

Note: Consider traffic load when assigning priority values to PGID 15. Heavy traffic in this group may restrict the bandwidth available to other groups.

Configuring ETS

Consider an example consistent with that used for port-based PFC configuration (on page 211):

Table 20. *ETS Configuration*

Priority	Usage	PGID	Bandwidth
0	LAN (best effort delivery)	0	10%
1	LAN (best effort delivery)		
2	LAN (best effort delivery)		
3	SAN (Fibre Channel over Ethernet, with PFC)	1	20%
4	Business Critical LAN (lossless Ethernet, with PFC)	2	30%
5	Latency-sensitive LAN	3	40%
6	Latency-sensitive LAN		
7	Network Management (strict)	15	unlimited

The example shown in Table 20 is only slightly different than the default configuration shown in Figure 19 on page 213. In this example, latency-sensitive LAN traffic (802.1p priority 5 through 6) are moved from priority group 2 to priority group 3. This leaves Business Critical LAN traffic (802.1p priority 4) in priority group 2 by itself. Also, a new group for network management traffic has been assigned. Finally, the bandwidth allocation for priority groups 1, 2, and 3 are revised.

Note: DCBX may be configured to permit sharing or learning PFC configuration with or from external devices. This example assumes that PFC configuration is being performed manually. See "Data Center Bridging Capability Exchange" on page 219 for more information on DCBX.

This example can be configured using the following commands:

1. Turn CEE on.

```
SI 4091(config)# cee enable
```

Note: Turning CEE on or off will automatically change some 802.1p QoS and 802.3x standard flow control settings (see "Turning CEE On or Off" on page 201).

2. Configure each allocated priority group with a description (optional), list of 802.1p priority values, and bandwidth allocation:

```
SI 4091(config)# cee global ets priority-group pgid 0 priority 0,1,2
                                        (Assign priorities for Regular LAN group)
SI 4091(config)# cee global ets priority-group pgid 0 desc "Regular LAN"
                                        (Set a group description—optional)
SI 4091(config)# cee global ets priority-group pgid 1 priority 3
                                        (Assign priorities for SAN group)
SI 4091(config)# cee global ets priority-group pgid 1 desc "SAN"
                                        (Set a group description—optional)
SI 4091(config)# cee global ets priority-group pgid 2 priority 4
                                        (Assign priorities for the critical group)
SI 4091(config)# cee global ets priority-group pgid 2 desc "Biz-Critical
LAN"
                                        (Set a group description—optional)
SI 4091(config)# cee global ets priority-group pgid 3 priority 5,6 gid 0
   bandwidth 10 pgid 1 bandwidth 20 pgid 2 bandwidth 30 pgid 3 bandwidth
   40
                                        (Configure the new priority group and
                                        set link bandwidth allocations)
SI 4091(config)# cee global ets priority-group pgid 3 desc
   "Latency-Sensitive LAN"
                                        (Set a group description—optional)
```

Notes:

- In this example, the command to configure the new priority group (PGID 3) is combined with the command to revise the link bandwidth for all groups. This is required to prevent PGID 3 from having a bandwidth allocation of 0% when initially created (which is not allowed).
- Because PGIDs 0, 1, and 2 existed as part of the prior default configuration, priority values may be redistributed among them using individual commands.

3. Configure the strict priority group with a description (optional) and a list of 802.1p priority values:

```
SI 4091(config)# cee global ets priority-group pgid 15 priority 7
                                        (Assign priority for the strict traffic group)
SI 4091(config)# cee global ets priority-group pgid 15 desc "Network
Management"
                                        (Set a group description—optional)
```

Note: Priority group 15 is permitted unlimited bandwidth. As such, the commands for priority group 15 do not include bandwidth allocation.

4. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

Data Center Bridging Capability Exchange

Data Center Bridging Capability Exchange (DCBX) protocol is a vital element of CEE. DCBX allows peer CEE devices to exchange information about their advanced capabilities. Using DCBX, neighboring network devices discover their peers, negotiate peer configurations, and detect misconfigurations.

DCBX provides two main functions on the SI4091:

- Peer information exchange

  The SI4091 uses DCBX to exchange information with connected CEE devices. For normal operation of any FCoE implementation on the SI4091, DCBX must remain enabled on all ports participating in FCoE.

- Peer configuration negotiation

  DCBX also allows CEE devices to negotiate with each other for the purpose of automatically configuring advanced CEE features such as PFC, ETS, and (for some CNAs) FIP. The administrator can determine which CEE feature settings on the SI4091 are communicated to and matched by CEE neighbors, and also which CEE feature settings on the SI4091 may be configured by neighbor requirements.

The DCBX feature requires CEE to be turned on (see "Turning CEE On or Off" on page 201).

DCBX Settings

When CEE is turned on, DCBX is enabled for peer information exchange on all ports. For configuration negotiation, the following default settings are configured:

- Application Protocol: FCoE and FIP snooping is set for traffic with 802.1p priority 3
- PFC: Enabled on 802.1p priority 3
- ETS
  - Priority group 0 includes priority values 0 through 2, with bandwidth allocation of 10%
  - Priority group 1 includes priority value 3, with bandwidth allocation of 50%
  - Priority group 2 includes priority values 4 through 7, with bandwidth allocation of 40%

Enabling and Disabling DCBX

When CEE is turned on, DCBX can be enabled and disabled on a per-port basis, using the following CLI command:

```
SI 4091(config)# [no] cee port <port alias or number> dcbx enable
```

When DCBX is enabled on a port, Link Layer Detection Protocol (LLDP) is used to exchange DCBX parameters between CEE peers. Also, the interval for LLDP transmission time is set to one second for the first five initial LLDP transmissions, after which it is returned to the administratively configured value. The minimum delay between consecutive LLDP frames is also set to one second as a DCBX default.

Peer Configuration Negotiation

CEE peer configuration negotiation can be set on a per-port basis for a number of CEE features. For each supported feature, the administrator can configure two independent flags:

- The `advertise` flag

  When this flag is set for a particular feature, the SI4091 settings will be transmit to the remote CEE peer. If the peer is capable of the feature, and willing to accept the SI4091 settings, it will be automatically reconfigured to match the SI4091.

- The `willing` flag

  Set this flag when required by the remote CEE peer for a particular feature as part of DCBX signaling and support. Although some devices may also expect this flag to indicate that the SI4091 will accept overrides on feature settings, the SI4091 retains its configured settings. As a result, the administrator should configure the feature settings on the SI4091 to match those expected by the remote CEE peer.

These flags are available for the following CEE features:

- Application Protocol

  DCBX exchanges information regarding FCoE and FIP snooping, including the 802.1p priority value used for FCoE traffic. The `advertise` flag is set or reset using the following command:

  ```
  SI 4091(config)# [no] cee port <port alias or number> dcbx app_proto
  advertise
  ```

  The `willing` flag is set or reset using the following command:

  ```
  SI 4091(config)# [no] cee port <port alias or number> dcbx app_proto willing
  ```

- PFC

  DCBX exchanges information regarding whether PFC is enabled or disabled on the port. The `advertise` flag is set or reset using the following command:

```
SI 4091(config)# [no] cee port <port alias or number> dcbx pfc advertise
```

  The willing flag is set or reset using the following command:

```
SI 4091(config)# [no] cee port <port alias or number> dcbx pfc willing
```

- ETS

  DCBX exchanges information regarding ETS priority groups, including their 802.1p priority members and bandwidth allocation percentages. The `advertise` flag is set or reset using the following command:

```
SI 4091(config)# [no] cee port <port alias or number> dcbx ets advertise
```

  The `willing` flag is set or reset using the following command:

```
SI 4091(config)# [no] cee port <port alias or number> dcbx ets willing
```

Configuring DCBX

Consider an example consistent Figure on page 198 and used with the previous FCoE examples in this chapter:

- FCoE is used on ports INT1 and EXT4.
- CEE features are also used with LANs on ports INT2 and EXT5.
- All other ports are disabled or are connected to regular (non-CEE) LAN devices.

In this example, the SI4091 acts as the central point for CEE configuration. FCoE-related ports will be configured for advertising CEE capabilities, but not to accept external configuration. Other LAN ports that use CEE features will also be configured to advertise feature settings to remote peers, but not to accept external configuration. DCBX will be disabled on all non-CEE ports.

This example can be configured using the following commands:

1. Turn CEE on.

```
SI 4091(config)# cee enable
```

Note: Turning CEE on or off will automatically change some 802.1p QoS and 802.3x standard flow control settings (see "Turning CEE On or Off" on page 201).

2. Enable desired DCBX configuration negotiation on FCoE ports:

```
SI 4091(config)# cee port INT1 dcbx enable
SI 4091(config)# cee port INT1 dcbx app_proto adv
SI 4091(config)# cee port INT1 dcbx pfc adv
SI 4091(config)# cee port INT1 dcbx ets adv

SI 4091(config)# cee port EXT4 dcbx enable
SI 4091(config)# cee port EXT4 dcbx app_proto adv
SI 4091(config)# cee port EXT4 dcbx pfc adv
SI 4091(config)# cee port EXT4 dcbx ets adv
```

To verify DCBX settings for each port, use the following command:

```
SI 4091(config)# show cee port INT1 dcbx
Current DCBX Configuration on port(s) INT1:

Alias Port DCBX State Feature     State     Willing   Advertise
===== ==== ========== =========== ========= ========= =========
INT1  1    Enabled    ETS         Enabled   Disabled  Enabled
INT1  1    Enabled    PFC         Enabled   Disabled  Enabled
INT1  1    Enabled    App Proto   Enabled   Disabled  Enabled
```

3. Enable desired DCBX advertisements on other CEE ports:

```
SI 4091(config)# cee port INT2 dcbx enable
SI 4091(config)# cee port INT2 dcbx pfc adv
SI 4091(config)# cee port INT2 dcbx ets adv

SI 4091(config)# cee port EXT5 dcbx enable
SI 4091(config)# cee port EXT5 dcbx pfc adv
SI 4091(config)# cee port EXT5 dcbx ets adv
```

4. Disable DCBX for each non-CEE port as appropriate:

```
SI 4091(config)# no cee port INT3 dcbx enable
SI 4091(config)# no cee port INT4 dcbx enable
...
SI 4091(config)# no cee port EXT10 dcbx enable
```

5. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

FCoE Example Configuration

The following example collects the various components from previous sections of this chapter.

Figure 20. A Mixed Fibre Channel and FCoE Network

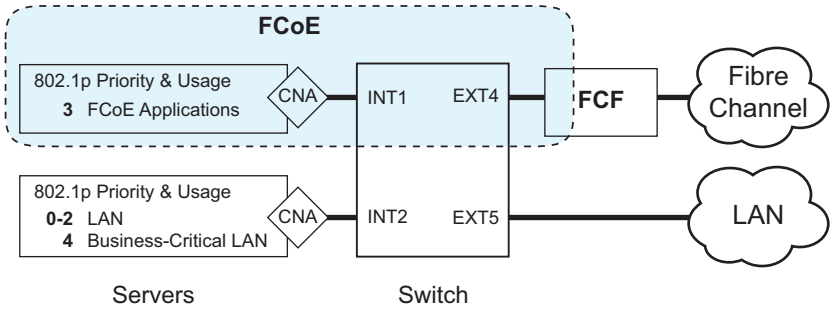

In Figure 20 on page 223, the Fibre Channel network is connected to the FCoE network through an FCF bridge module on port EXT4. The FCoE-enabled SI4091 is internally connected to a blade server (ENode) through an FCoE-enabled CNA on port INT1.

1. Turn CEE on:

```
SI 4091(config)# cee enable
```

Note: Turning CEE on or off will automatically change some 802.1p QoS and 802.3x standard flow control settings (see "Turning CEE On or Off" on page 201).

2. Disable FIP snooping on all non-FCoE external ports:

```
SI 4091(config)# no fcoe fips port ext1-ext3,ext5-ext10 enable
```

3. Turn global FIP snooping on:

```
SI 4091(config)# fcoe fips enable
```

4. Configure the FCoE VLAN:

```
SI 4091(config)# vlan 1002              (Select and enable the VLAN)
SI 4091(config-vlan)# exit              (Exit VLAN configuration mode)
```

5. Configure the FCoE ports and enable VLAN tagging.

   Place all FCoE ports associated with each FCF into a unique VLAN supported by that FCF:

```
SI 4091(config)# interface port INT1,EXT4          (Select FCoE ports)
SI 4091(config-if)# switchport mode trunk          (Enable VLAN tagging)
SI 4091(config-if)# switchport trunk allowed vlan 1,1002
SI 4091(config-if)# exit                           (Exit port configuration mode)
```

Note: Placing ports into the VLAN *after* tagging is enabled helps to ensure that their port VLAN ID (PVID) is not accidentally changed.

6. For each FCoE port, enable FIP snooping if necessary (set by default), and set the desired FCF mode:

```
SI 4091(config)# fcoe fips port INT1,EXT4 enable (Enable FIPS on FCoE ports)
SI 4091(config)# fcoe fips port INT1 fcf-mode off(Set as ENode connection)
SI 4091(config)# fcoe fips port EXT4 fcf-mode on (Set as FCF connection)
```

7. Enable PFC for the FCoE traffic, if necessary.

   Note: By default, PFC is enabled on priority 3 for all ports. If using the default settings, the commands in this step are not necessary and are shown here merely as an example of manual configuration.

```
SI 4091(config)# cee port INT1 pfc      (Turn on PFC for port INT1)
SI 4091(config)# cee port INT1 pfc priority 3 enable
                                        (Select and enable FCoE priority 3 PFC for the port)
SI 4091(config)# cee port INT1 pfc priority 3 description "FCoE"
                                        (Set priority description—optional)

SI 4091(config)# cee port EXT4 pfc      (Select and enable FCoE for port EXT4)
SI 4091(config)# cee port EXT4 pfc priority 3 enable
                                        (Select and enable FCoE priority 3 PFC for the port)
SI 4091(config)# cee port EXT4 pfc priority 3 description "FCoE"
                                        (Set priority description—optional)
```

8. Enable PFC for the business-critical LAN application:

```
SI 4091(config)# cee port INT2 pfc      (Turn on PFC for port INT2)
SI 4091(config)# cee port INT2 pfc priority 4 enable
                                        (Select and enable FCoE priority 4 PFC for the port)
SI 4091(config)# cee port INT2 pfc priority 4 description
"Business-critical LAN"
                                        (Set priority description—optional)
SI 4091(config)# cee port EXT5 pfc      (Select and enable FCoE for port EXT5)
SI 4091(config)# cee port EXT5 pfc priority 4 enable
                                        (Select and enable FCoE priority 4 PFC for the port)
SI 4091(config)# cee port EXT5 pfc priority 4 description
"Business-critical LAN"
                                        (Set priority description—optional)
```

Notes:

- By default, PFC is enabled for all ports. If the default settings are in place, the two `cee port` <*number*> `pcf` commands may be omitted. The commands are shown here as a manual configuration example only.
- Although PFC is enabled on the ports by default, it is not enabled on priority 4 by default. The `cee port` <*number*> `pfc priority` <*PGID*> `enable` commands are still required.

9. Configure each allocated priority group with a description (optional), list of 802.1p priority values, and bandwidth allocation:

```
SI 4091(config)# cee global ets priority-group pgid 0 priority 0,1,2
                                        (Assign priorities for Regular LAN group)
SI 4091(config)# cee global ets priority-group pgid 0 desc "Regular LAN"
                                        (Set a group description—optional)

SI 4091(config)# cee global ets priority-group pgid 1 priority 3
                                        (Assign priorities for SAN group)
SI 4091(config)# cee global ets priority-group pgid 1 desc "SAN"
                                        (Set a group description—optional)

SI 4091(config)# cee global ets priority-group pgid 2 priority 4
                                        (Assign priorities for the critical group)
SI 4091(config)# cee global ets priority-group pgid 2 desc "Biz-Critical
LAN"
                                        (Set a group description—optional)

SI 4091(config)# cee global ets priority-group pgid 3 priority 5,6 gid 0
   bandwidth 10 pgid 1 bandwidth 20 pgid 2 bandwidth 30 pgid 3 bandwidth
   40
                                        (Configure the new priority group and
                                        set link bandwidth allocations)
SI 4091(config)# cee global ets priority-group pgid 3 desc
   "Latency-Sensitive LAN"
                                        (Set a group description—optional)
```

Notes:

- In this example, the command to configure the new priority group (PGID 3) is combined with the command to revise the link bandwidth for all groups. This is required to prevent PGID 3 from having a bandwidth allocation of 0% when initially created (which is not allowed).
- Because PGIDs 0, 1, and 2 existed as part of the prior default configuration, priority values may be redistributed among them using individual commands.
- Do not place PFC-enabled traffic queues and non-PFC traffic queues in the same priority group, as this may result in unexpected network behavior.

10. Configure the strict priority group with a description (optional) and a list of 802.1p priority values:

```
SI 4091(config)# cee global ets priority-group pgid 15 priority 7
                                        (Assign priority for the strict traffic group)
SI 4091(config)# cee global ets priority-group pgid 15 desc "Network
Management"
                                        (Set a group description—optional)
```

Note: Priority group 15 is permitted unlimited bandwidth. As such, the commands for priority group 15 do not include bandwidth allocation.

11. Enable desired DCBX configuration negotiation on FCoE ports:

```
SI 4091(config)# cee port INT1 dcbx enable
SI 4091(config)# cee port INT1 dcbx app_proto adv
SI 4091(config)# cee port INT1 dcbx pfc adv
SI 4091(config)# cee port INT1 dcbx ets adv

SI 4091(config)# cee port EXT4 dcbx enable
SI 4091(config)# cee port EXT4 dcbx app_proto adv
SI 4091(config)# cee port EXT4 dcbx pfc adv
SI 4091(config)# cee port EXT4 dcbx ets adv
```

12. Enable desired DCBX advertisements on other CEE ports:

```
SI 4091(config)# cee port INT2 dcbx enable
SI 4091(config)# cee port INT2 dcbx pfc adv
SI 4091(config)# cee port INT2 dcbx ets adv

SI 4091(config)# cee port EXT5 dcbx enable
SI 4091(config)# cee port EXT5 dcbx pfc adv
SI 4091(config)# cee port EXT5 dcbx ets adv
```

13. Disable DCBX for each non-CEE port as appropriate:

```
SI 4091(config)# no cee port INT3 dcbx enable
SI 4091(config)# no cee port INT4 dcbx enable
...
SI 4091(config)# no cee port EXT10 dcbx enable
```

14. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

Part 6: High Availability Fundamentals

Internet traffic consists of myriad services and applications which use the Internet Protocol (IP) for data delivery. However, IP is not optimized for all the various applications. High Availability goes beyond IP and makes intelligent switching decisions to provide redundant network configurations.

Chapter 17. Layer 2 Failover

The primary application for Layer 2 Failover is to support Network Adapter Teaming. With Network Adapter Teaming, all the NICs on each server share the same IP address, and are configured into a team. One NIC is the primary link, and the other is a standby link. For more details, refer to the documentation for your Ethernet adapter.

Note: Only two links per server blade can be used for Layer 2 LAG Failover (one primary and one backup). Network Adapter Teaming allows only one backup NIC for each server blade.

There are two types of Layer 2 Failover:

- Manual Monitoring (MMON)

  MMON allows you to specify a set of ports and/or LAGs to be monitored for link health. If a configurable number of monitored links fails, all ports specified in a control set are disabled in order to trigger NIC failover.

  MMON is enabled by default.

- Automatic Monitoring (AMON)

  When AMON is enabled on any LAG, if a configurable number of member links fails, all internal ports (or those for affected VLANs if VLAN monitoring is active) are disabled to trigger NIC failover.

Note: MMON and AMON failover are mutually exclusive. They cannot both be configured to operate on the SI4091 at the same time.

Manual Monitoring Aggregation Links

MMON allows you to configure a set of external ports and/or LAGs to be monitored for link failures (a monitor list), and another set of internal and/or external ports or LAGs to disable when a failure trigger limit is reached (a control list). When the SI4091 detects a certain number of link failures (configurable) in the monitor list, it automatically disables all items in the control list.

Consider a scenario where a specific set of servers (attached to SI4091 internal ports) relies upon not just one external uplink port, but on the aggregate bandwidth available in a multi-port LAG. Without Layer 2 Failover, if one or more of the external ports in the LAG were to fail, the overall service capacity might become degraded even though a significant portion of the LAG remains available. In such a case, it may be preferable to force the affected servers to failover to their alternate NIC.

Layer 2 Failover helps to accomplish NIC failover. In the above scenario, the external LAG is placed on the monitor list, and the corresponding internal (server) ports are placed on the control list. Then, if a partial LAG failure occurred, the server's active network adapters would detect the disabled internal link and trigger a network-adapter failover to an alternate port or LAG on the SI4091, or an alternate SI4091 in the chassis.

The SI4091 automatically enables the control list items when the monitor list items return to service.

MMON Default Settings

By default, Layer 2 Failover is enabled and pre-configured for MMON operation with Trigger 1:

- Monitor list: EXT1–EXT10 (in LAG 25 with LACP admin key 1000)
- Control list: Internal ports INT1–INT14

By default, one Layer 2 Failover trigger is configured. All external ports are placed on the monitor list, and all internal ports are placed on the control list. Because the default failover limit for each trigger is 0, when there no external ports are connected, all internal ports are automatically disabled.

MMON Port States

Monitor Port State

A monitor port is considered operational as long as the following conditions are true:

- The port must be in the `Link Up` state.
- If the port is part of an LACP LAG, the port must be in the `Aggregated` state.

If any of the above conditions is false, the monitor port is considered to have failed.

Only external ports and external LAGs can be placed on the monitor list.

Control Port State

A control port is considered Operational if the monitor trigger is up. As long as the trigger is up, the port is considered operational from a teaming perspective, even if the port itself is actually in the `Down` state or `Not Aggregated` state (if part of an LACP LAG).

A control port is considered to have failed only if the monitor trigger is in the `Down` state.

Internal and external ports and LAGs can be placed on the control list.

Viewing the Port Status

To view the state of any port, use one of the following commands:

```
>> # show interface link              (View port link status)
>> # show lacp information            (View port LACP status)
```

MMON Failover Limits

The failover limit lets you specify the minimum number of operational links required within each trigger before the trigger initiates a failover event. For example, if the limit is two, a failover event occurs when the number of operational links in the trigger is two or fewer.

When you set the limit to zero (the default for each trigger), the SI4091 initiates a failover event only when no links in the trigger are operational.

MMON with LACP

L2 Failover works with Link Aggregation Control Protocol (LACP) as follows.

Link Aggregation Control Protocol allows the switch to form LAGs. You can use the *admin key* to add up to two LACP LAGs to a failover trigger using manual monitoring. When you add an *admin key* to a trigger, any LACP LAG with that *admin key* becomes a member of the trigger.

MMON Configuration Guidelines

This section provides important information about configuring Layer 2 Failover with MMON.

- MMON and AMON failover are mutually exclusive. They cannot both be configured to operate on the SI4091 at the same time.
- MMON can monitor only external ports. Internal ports cannot be placed on the monitor list.
- Any specific failover trigger can monitor external ports only, static LAGs only, or LACP LAGs only. The different types cannot be combined in the same trigger.
- Port membership for different triggers should not overlap. Any specific port should be a member of only one trigger.

Configuring MMON

Use the following procedure to configure a Layer 2 Failover Manual Monitor.

1. Specify the links to monitor.

```
SI 4091(config)# failover trigger 1 mmon monitor member ext1-ext5
```

2. Specify the links to disable when the failover limit is reached.

```
SI 4091(config)# failover trigger 1 mmon control member int1-int8
```

3. Configure general Failover parameters.

```
SI 4091(config)# failover enable          (Enable Layer 2 Failover)
SI 4091(config)# failover trigger 1 enable(Enable the individual trigger)
SI 4091(config)# failover trigger 1 limit 2(Set the failover limit)
```

4. Verify the configuration.

```
SI 4091(config)# show failover trigger 1 information
```

Automatic Monitoring LAG Links

AMON can be enabled on any LAG in the SI4091, including LACP LAGs. LAGs can be added to failover trigger groups. Then, if some specified number of trigger links fail, the SI4091 disables all the internal ports in the SI4091 (unless VLAN Monitor is turned on). When the internal ports are disabled, it causes the NIC team on the affected server blades to failover from the primary to the backup NIC. This process is called a failover event.

When the appropriate number of links in a trigger group return to service, the SI4091 enables the internal ports. This causes the NIC team on the affected server blades to fail back to the primary SI4091 (unless Auto-Fallback is disabled on the NIC team). The backup SI4091 processes traffic until the primary SI4091's internal links come up, which can take up to five seconds.

VLAN Monitoring

The VLAN Monitor allows Layer 2 Failover to discern different VLANs. With VLAN Monitor turned on:

- If enough links in a trigger fail (see "MMON Failover Limits" on page 231), the SI4091 disables all internal ports that reside in the same VLAN membership as the LAG(s) in the trigger.
- When enough links in the trigger return to service, the SI4091 enables the internal ports that reside in the same VLAN membership as the LAG(s) in the trigger.

If you turn off the VLAN Monitor (`SI 4091# no failover vlan`), only one failover trigger is allowed. When a link failure occurs on the trigger, the SI4091 disables all internal server-blade ports.

AMON Topologies

Figure 21 is a simple example of Layer 2 Failover. One SI4091 is the primary, and the other is used as a backup. In this example, all external ports on the primary SI4091 belong to a single LAG, with Layer 2 Failover enabled, and Failover Limit set to 2. If two or fewer links in trigger 1 remain active, the SI4091 temporarily disables all internal server-blade ports that reside in VLAN 1. This action causes a failover event on Server 1 and Server 2.

Figure 21. Basic Layer 2 Failover

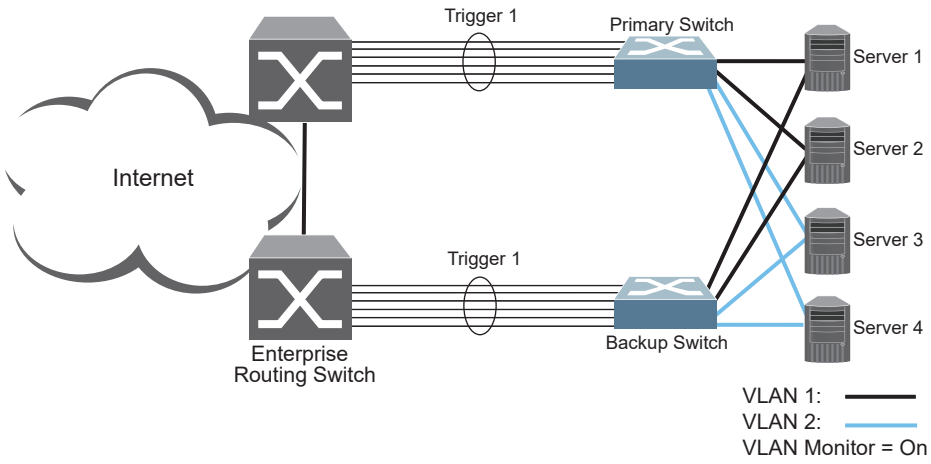

Figure 22 shows a configuration with two LAGs, each in a different Failover Trigger. SI4091 #1 is the primary for Server 1 and Server 2. SI4091 #2 is the primary for Server 3 and Server 4. VLAN Monitor is turned on.

If all links go down in trigger 1, SI4091 #1 disables all internal ports that reside in VLAN 1. If all links in trigger 2 go down, SI4091 #1 disables all internal ports that reside in VLAN 2.

Figure 22. Two LAGs, each in a different Failover Trigger

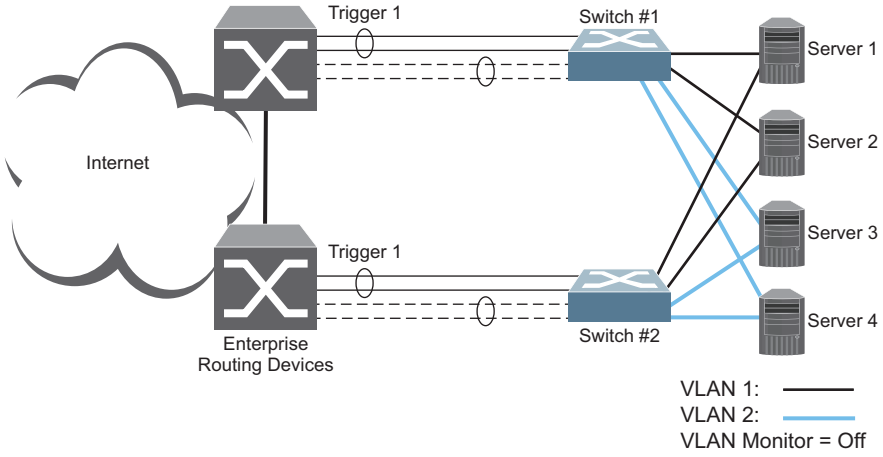

Figure 23 shows a configuration with two LAGs. VLAN Monitor is turned off, so only one Failover Trigger is configured on each SI4091. SI4091 #1 is the primary for Server 1 and Server 2. SI4091 #2 is the primary for Server 3 and Server 4..

If all links in trigger 1 go down, SI4091 #1 disables all internal links to server blades.

Figure 23. Two LAGs, one Failover Trigger

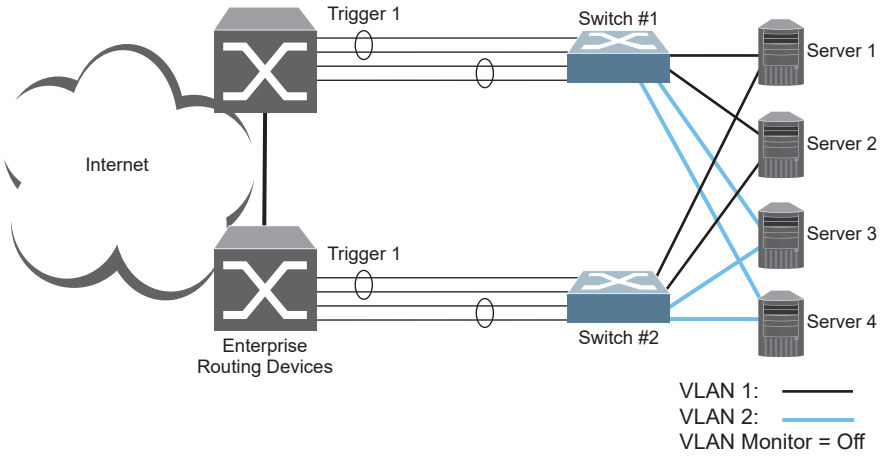

AMON Failover Limit

The failover limit lets you specify the minimum number of operational links required within each trigger before the trigger initiates a failover event. For example, if the limit is two, a failover event occurs when the number of operational links in the trigger is two or fewer. When you set the limit to zero, the SI4091 triggers a failover event only when no links in the trigger are operational.

AMON with LACP

AMON works with Link Aggregation Control Protocol (LACP) as follows.

Link Aggregation Control Protocol allows the switch to form dynamic LAGs. You can use the *admin key* to add LACP LAGs to a failover trigger using automatic monitoring. When you add an *admin key* to a trigger, any LACP LAG with that *admin key* becomes a member of the trigger.

AMON Configuration Guidelines

This section provides important information about configuring Layer 2 Failover with AMON.

- MMON and AMON failover are mutually exclusive. They cannot both be configured to operate on the SI4091 at the same time.
- Any specific failover trigger may monitor static LAGs only or LACP LAGs only, but not both.
- All external ports in all static or LACP LAGs added to any specific failover trigger must belong to the same VLAN.
- When VLAN Monitor is on, the following additional guidelines apply:
  - All external ports in all static or LACP LAGs added to a specific failover trigger must belong to the same VLAN and have the same PVID.
  - Different triggers are not permitted to operate on the same VLAN.
  - Different triggers are not permitted to operate on the same internal port.

AMON Configuration Example

The following procedure pertains to the configuration shown in Figure 21.

1. Configure Network Adapter Teaming on the servers.
2. Define a LAG on the SI4091.

```
SI 4091(config)# portchannel 1 port EXT1,EXT2,EXT3 enable
```

3. Configure Failover trigger parameters.

```
SI 4091(config)# failover trigger 1 enable
SI 4091(config)# failover trigger 1 limit 2
SI 4091(config)# failover trigger 1 amon portchannel 1
```

4. Configure general Failover parameters.

```
SI 4091(config)# failover enable
```

5. Verify the configuration.

```
SI 4091(config)# show failover trigger 1 information
```

Chapter 18. Hot Links

The following topics are discussed in this chapter:

Hot Links Overview

Hot Links provides basic link redundancy with fast recovery.

Up to 25 Hot Links triggers can be configured. A basic trigger consists of a pair of layer 2 interfaces, each containing an individual port, LAG, or LACP adminkey. One interface is the Master, and the other is a Backup. While the Master interface is in the active state (forwarding traffic), the Backup interface is set to the standby state (blocks traffic). In the event that the Master fails, the Backup will transition to the active state and forward traffic.

By default, Hot Links favors link stability. To minimize disruption after failover, the Master interface does not automatically resume active status once it is restored (unless the preemption option is enabled). Instead, the recovered Master transitions to the standby state and blocks traffic while the Backup continues as the active link. However, if the Backup fails, the Master will return to active status.

You may select any external port, static LAG, or an LACP adminkey as part of a Hot Link pair. Internal ports cannot participate in Hot Links triggers.

Hot Link triggers can also be associated with specific sets of VLANs. If a trigger is configured with a set of VLANs, only those VLANs will transition to the new active link at failover.

Hot Links also offers VLAN load-balancing. When enabled on a trigger, the Master and Backup interfaces are configured such that both are active for half of the associated VLANs: half on the Master interface and half on the Backup. However if either interface fails, its partner will assume the active role for the full set of configured VLANs.

Hot Links Options

The following options are available with Hot LInks.

Forward Delay

The Forward Delay timer allows Hot Links to monitor the Master and Backup interfaces for link stability before transition one to the active state. Before the transition occurs, the interface must maintain a stable link for the duration of the Forward Delay interval.

For example, consider a Forward Delay timer set to 10 seconds for trigger 1:

```
SI 4091(config)# hotlinks trigger 1 forward-delay 10
```

In this case, the SI4091 will select an interface to become active only if its link remained stable for 10 seconds. Otherwise, the link is considered unstable and the Forward Delay timer restarts.

Preemption

If you prefer that the Master interface returns to active state after recovering from a failure, preempting the Backup, enable the Hot Links preemption option.

When preemption is enabled, the Master interface transitions to the active state immediately upon recovery and the Backup interface is immediately forced to transition to the standby state. If Forward Delay is enabled, the preemption transition occurs only once the Master interface has maintained link stability for the duration of the Forward Delay period.

FDB Update

Use the FDB update option to notify other devices on the network about updates to the Forwarding Database (FDB). When you enable FDB update, the SI4091 sends multicasts of addresses in the forwarding database (FDB) over the active interface so that other devices on the network can learn the new path. The Hot Links FBD update option uses the station update rate to determine the rate at which to send FDB packets.

Note: If you clear the FDB table at least twice while traffic is running with less than two seconds between the clear FDB commands, some flooded unicast packets will appear in the software FDB but not in the hardware table. To avoid this, wait at least three seconds before reissuing a command to clear the FDB table.

Configuration Guidelines

The following configuration guidelines apply to Hot links:

- Only external ports, LAGs, and LACP adminkeys can be configured as Hot Links. Internal ports cannot participate in Hot Links.
- A port that is a member of the Master interface cannot be a member of the Backup interface. A port that is a member of one Hot Links trigger cannot be a member of any other Hot Links trigger.
- An individual port that is configured as a Hot Link interface cannot be added as a member of a LAG.

Configuring Hot Links

Hot Links can be configured for a variety of strategies, as outlined in the examples in this section. One configured, Hot Link triggers can be used as standard interface entities, just like individual ports or LAGs.

Example 1: Port-Based Hot Links

In the following example, Hot Links behavior is based purely on port. VLAN-based load-balancing is disabled. EXT1 is the Master interface, regardless of VLAN. EXT2 is the Backup and remains in standby mode unless EXT1 fails.

```
SI 4091(config)# hotlinks trigger 1 enable (Enable Hot Links Trigger 1)
SI 4091(config)# hotlinks trigger 1 master port ext1 (Add port to Master interface)
SI 4091(config)# hotlinks trigger 1 backup port ext2(Add port to Backup interface)
SI 4091(config)# no hotlinks trigger 1 backup prefer(Use port-based Hot Links)
SI 4091(config)# hotlinks enable       (Turn on Hot Links)
```

Example 2: Automatic VLAN Load-Balancing

In the following example, Hot Links are configured with automatic VLAN load-balancing. EXT1 is active for half of the VLANs (selected automatically by the SI4091). EXT2 is active for the other half. In the event that either interface fails, the remaining interface will become active for the full set of VLANs.

```
SI 4091(config)# hotlinks trigger 1 enable (Enable Hot Links Trigger 1)
SI 4091(config)# hotlinks trigger 1 master port ext1 (Add port to Master interface)
SI 4091(config)# hotlinks trigger 1 backup port ext2(Add port to Backup interface)
SI 4091(config)# hotlinks trigger 1 backup prefer auto(Automatic VLAN load-balancing)
SI 4091(config)# hotlinks enable       (Turn on Hot Links)
```

When automatic load-balancing is enabled, removing some portion of the VLANs (by deleting the VLAN from the interface) may result in an imbalance in the distribution already established among the interfaces.

After making changes to the VLAN configuration, manually rebalance the existing VLAN load among the participating Hot Links interface using the following command:

```
SI 4091(config)# hotlinks trigger 1 vlan rebalance
```

Example 3: VLAN Preference

In the following example, Hot Links are configured such that traffic on VLANs 50–60 prefers EXT1 while active. Other VLANs will be distributed to EXT1 or EXT2 as necessary for load-balancing. In the event that either interface fails, the remaining interface will become active for the full set of VLANs.

```
SI 4091(config)# hotlinks trigger 1 enable(Enable Hot Links Trigger 1)
SI 4091(config)# hotlinks trigger 1 master port ext1(Add port to Master interface)
SI 4091(config)# hotlinks trigger 1 backup port ext2(Add port to Backup interface)
SI 4091(config)# hotlinks trigger 1 backup prefer vlan 50-60(Set VLAN participation)
SI 4091(config)# hotlinks enable       (Turn on Hot Links)
```

Part 7: Network Management

Chapter 19. Link Layer Discovery Protocol

The SI4091 supports Link Layer Discovery Protocol (LLDP). This chapter discusses the use and configuration of LLDP on the switch:

LLDP Overview

Link Layer Discovery Protocol (LLDP) is an IEEE 802.1AB-2005 standard for discovering and managing network devices. LLDP uses Layer 2 (the data link layer), and allows network management applications to extend their awareness of the network by discovering devices that are direct neighbors of already known devices.

With LLDP, the SI4091 can advertise the presence of its ports, their major capabilities, and their current status to other LLDP stations in the same LAN. LLDP transmissions occur on ports at regular intervals or whenever there is a relevant change to their status. The switch can also receive LLDP information advertised from adjacent LLDP-capable network devices.

In addition to discovery of network resources, and notification of network changes, LLDP can help administrators quickly recognize a variety of common network configuration problems, such as unintended VLAN exclusions or mis-matched port aggregation membership.

The LLDP transmit function and receive function can be independently configured on a per-port basis. The administrator can allow any given port to transmit only, receive only, or both transmit and receive LLDP information.

The LLDP information to be distributed by the SI4091 ports, and that which has been collected from other LLDP stations, is stored in the switch's Management Information Base (MIB). Network Management Systems (NMS) can use Simple Network Management Protocol (SNMP) to access this MIB information. LLDP-related MIB information is read-only.

Changes, either to the local switch LLDP information or to the remotely received LLDP information, are flagged within the MIB for convenient tracking by SNMP-based management systems.

For LLDP to provide expected benefits, all network devices that support LLDP should be consistent in their LLDP configuration.

Enabling or Disabling LLDP

Follow the instructions to enable or disable LLDP.

Global LLDP Setting

By default, LLDP is enabled on the SI4091. To turn LLDP off or on, use the following commands:

```
SI 4091(config)# [no] lldp enable       (Turn LLDP on or off globally)
```

Transmit and Receive Control

The SI4091 can also be configured to transmit or receive LLDP information on a port-by-port basis. By default, when LLDP is globally enabled on the switch, SI4091 ports transmit and receive LLDP information (see the `tx_rx` option below). To change the LLDP transmit and receive state, the following commands are available:

```
SI 4091(config)# interface port <n>      (Select a switch port)
SI 4091(config-if)# lldp admin-status tx_rx(Transmit and receive LLDP)
SI 4091(config-if)# lldp admin-status tx_only(Only transmit LLDP)
SI 4091(config-if)# lldp admin-status rx_only(Only receive LLDP)
SI 4091(config-if)# no lldp admin-status(Do not participate in LLDP)
SI 4091(config-if)# exit                 (Exit port mode)
```

To view the LLDP transmit and receive status, use the following commands:

```
SI 4091(config)# show lldp port          (status of all ports)
SI 4091(config)# show interface port <n> lldp(status of selected port)
```

LLDP Transmit Features

Numerous LLDP transmit options are available, including scheduled and minimum transmit interval, expiration on remote systems, SNMP trap notification, and the types of information permitted to be shared.

Scheduled Interval

The SI4091 can be configured to transmit LLDP information to neighboring devices once each 5 to 32768 seconds. The scheduled interval is global; the same interval value applies to all LLDP transmit-enabled ports. However, to help balance LLDP transmissions and keep them from being sent simultaneously on all ports, each port maintains its own interval clock, based on its own initialization or reset time. This allows switch-wide LLDP transmissions to be spread out over time, though individual ports comply with the configured interval.

The global transmit interval can be configured using the following command:

```
SI 4091(config)# lldp refresh-interval <interval>
```

where *interval* is the number of seconds between LLDP transmissions. The range is 5 to 32768. The default is 30 seconds.

Minimum Interval

In addition to sending LLDP information at scheduled intervals, LLDP information is also sent when the SI4091 detects relevant changes to its configuration or status (such as when ports are enabled or disabled). To prevent the SI4091 from sending multiple LLDP packets in rapid succession when port status is in flux, a transmit delay timer can be configured.

The transmit delay timer represents the minimum time permitted between successive LLDP transmissions on a port. Any interval-driven or change-driven updates will be consolidated until the configured transmit delay expires.

The minimum transmit interval can be configured using the following command:

```
SI 4091(config)# lldp transmission-delay <interval>
```

where *interval* is the minimum number of seconds permitted between successive LLDP transmissions on any port. The range is 1 to one-quarter of the scheduled transmit interval (`msgtxint`), up to 8192. The default is 2 seconds.

Time-to-Live for Transmitted Information

The transmitted LLDP information is held by remote systems for a limited time. A time-to-live parameter allows the switch to determine how long the transmitted data should be held before it expires. The hold time is configured as a multiple of the configured transmission interval:

```
SI 4091(config)# lldp holdtime-multiplier <multiplier>
```

where *multiplier* is a value between 2 and 10. The default value is 4, meaning that remote systems will hold the port's LLDP information for 4 **x** the 30-second `msgtxint` value, or 120 seconds, before removing it from their MIB.

Trap Notifications

If SNMP is enabled on the SI4091 (see "Using Simple Network Management Protocol" on page 30), each port can be configured to send SNMP trap notifications whenever LLDP transmissions are sent. By default, trap notification is disabled for each port. The trap notification state can be changed using the following commands):

```
SI 4091(config)# interface port 1
SI 4091(config-if)# [no] lldp trap-notification
SI 4091(config-if)# exit
```

In addition to sending LLDP information at scheduled intervals, LLDP information is also sent when the SI4091 detects relevant changes to its configuration or status (such as when ports are enabled or disabled). To prevent the SI4091 from sending multiple trap notifications in rapid succession when port status is in flux, a global trap delay timer can be configured.

The trap delay timer represents the minimum time permitted between successive trap notifications on any port. Any interval-driven or change-driven trap notices from the port will be consolidated until the configured trap delay expires.

The minimum trap notification interval can be configured using the following command:

```
SI 4091(config)# lldp trap-notification-interval <interval>
```

where *interval* is the minimum number of seconds permitted between successive LLDP transmissions on any port. The range is 1 to 3600. The default is 5 seconds.

If SNMP trap notification is enabled, the notification messages can also appear in the system log. This is enabled by default. To change whether the SNMP trap notifications for LLDP events appear in the system log, use the following command

```
SI 4091(config)# [no] logging log lldp
```

Changing the LLDP Transmit State

When the port is disabled, or when LLDP transmit is turned off for the port using the `admstat` command's `rx_only` or `disabled` options (see "Transmit and Receive Control" on page 247), a final LLDP packet is transmitted with a time-to-live value of 0. Neighbors that receive this packet will remove the LLDP information associated with the SI4091 port from their MIB.

In addition, if LLDP is fully disabled on a port (using `admstat disabled`) and later re-enabled, the SI4091 will temporarily delay resuming LLDP transmissions on the port in order to allow the port LLDP information to stabilize. The reinitialization delay interval can be globally configured for all ports using the following command::

```
SI 4091(config)# lldp reinit-delay <interval>
```

where *interval* is the number of seconds to wait before resuming LLDP transmissions. The range is between 1 and 10. The default is 2 seconds.

Types of Information Transmitted

When LLDP transmission is permitted on the port (see "Enabling or Disabling LLDP" on page 247), the port advertises the following required information in type/length/value (TLV) format:

- Chassis ID
- Port ID
- LLDP Time-to-Live

LLDP transmissions can also be configured to enable or disable inclusion of optional information, using the following command:

```
SI 4091(config)# interface port 1
SI 4091(config-if)# [no] lldp tlv <type>
SI 4091(config-if)# exit
```

where *type* is an LLDP information option from Table 21:

Table 21. *LLDP Optional Information Types*

Type	Description	Default
`portdesc`	Port Description	Enabled
`sysname`	System Name	Enabled
`sysdescr`	System Description	Enabled
`syscap`	System Capabilities	Enabled
`mgmtaddr`	Management Address	Enabled
`portvid`	IEEE 802.1 Port VLAN ID	Disabled
`portprot`	IEEE 802.1 Port and Protocol VLAN ID	Disabled
`vlanname`	IEEE 802.1 VLAN Name	Disabled

Table 21. *LLDP Optional Information Types (continued)*

Type	Description	Default
`protid`	IEEE 802.1 Protocol Identity	Disabled
`macphy`	IEEE 802.3 MAC/PHY Configuration/Status, including the auto-negotiation, duplex, and speed status of the port.	Disabled
`powermdi`	IEEE 802.3 Power via MDI, indicating the capabilities and status of devices that require or provide power over twisted-pair copper links.	Disabled
`linkaggr`	IEEE 802.3 Link Aggregation status for the port.	Disabled
`framesz`	IEEE 802.3 Maximum Frame Size for the port.	Disabled
`dcbx`	Data Center Bridging Capability Exchange Protocol (DCBX) for the port.	Enabled
`all`	Select all optional LLDP information for inclusion or exclusion.	Disabled

LLDP Receive Features

Types of Information Received

When the LLDP receive option is enabled on a port (see "Enabling or Disabling LLDP" on page 247), the port may receive the following information from LLDP-capable remote systems:

- Chassis Information
- Port Information
- LLDP Time-to-Live
- Port Description
- System Name
- System Description
- System Capabilities Supported/Enabled
- Remote Management Address

The SI4091 stores the collected LLDP information in the MIB. Each remote LLDP-capable device is responsible for transmitting regular LLDP updates. If the received updates contain LLDP information changes (to port state, configuration, LLDP MIB structures, deletion), the switch will set a change flag within the MIB for convenient notification to SNMP-based management systems.

Viewing Remote Device Information

LLDP information collected from neighboring systems can be viewed in numerous ways:

- Using a centrally-connected LLDP analysis server
- Using an SNMP agent to examine the SI4091 MIB
- Using CLI commands on the SI4091

Using the CLI, the following command displays remote LLDP information:

```
SI 4091(config)# show lldp remote-device [<index number>]
```

To view a summary of remote information, omit the *Index number* parameter. For example:

```
SI 4091(config)# show lldp remote-device
LLDP Remote Devices Information

LocalPort | Index | Remote Chassis ID    | Remote Port | Remote System
Name
----------|-------|---------------------|-------------|--------------
-------
3         | 1     | 00 18 b1 33 1d 00    | 23          |
```

To view detailed information for a remote device, specify the *Index number* as found in the summary. For example, in keeping with the sample summary, to list details for the first remote device (with an `Index` value of 1), use the following command:

```
SI 4091(config)# show lldp remote-device 1
Local Port Alias: 3
        Remote Device Index     : 1
        Remote Device TTL       : 99
        Remote Device RxChanges : false
        Chassis Type            : Mac Address
        Chassis Id              : 00-18-b1-33-1d-00
        Port Type               : Locally Assigned
        Port Id                 : 23
        Port Description        : 7

        System Name        :
        System Description : SI4091 System Interconnect Module,
                             flash image: version 6.9.0,
                             boot image: version 6.9.0

        System Capabilities Supported : bridge, router
        System Capabilities Enabled   : bridge, router

        Remote Management Address:
                Subtype             : IPv4
                Address             : 10.100.120.181
                Interface Subtype   : ifIndex
                Interface Number    : 128
                Object Identifier   :
```

Note: Received LLDP information can change very quickly. When using information commands, it is possible that flags for some expected events may be too short-lived to be observed in the output.

Time-to-Live for Received Information

Each remote device LLDP packet includes an expiration time. If the switch port does not receive an LLDP update from the remote device before the time-to-live clock expires, the switch will consider the remote information to be invalid, and will remove all associated information from the MIB.

Remote devices can also intentionally set their LLDP time-to-live to 0, indicating to the switch that the LLDP information is invalid and should be immediately removed.

LLDP Example Configuration

1. Turn LLDP on globally.

```
SI 4091(config)# lldp enable
```

2. Set the global LLDP timer features.

```
SI 4091(config)# lldp refresh-interval 30(Transmit each 30 seconds)
SI 4091(config)# lldp transmission-delay 2(No more often than 2 sec.)
SI 4091(config)# lldp holdtime-multiplier 4(Remote hold 4 intervals)
SI 4091(config)# lldp reinit-delay 2   (Wait 2 seconds after reinitialization)
SI 4091(config)# lldp trap-notification-interval 5(Minimum 5 sec. between)
```

3. Set LLDP options for each port.

```
SI 4091(config)# interface port <n>      (Select a switch port)
SI 4091(config-if)# lldp admin-status tx_rx(Transmit and receive LLDP)
SI 4091(config-if)# lldp trap-notification(Enable SNMP trap notifications)
SI 4091(config-if)# lldp tlv all         (Transmit all optional information)
SI 4091(config-if)# exit
```

4. Enable syslog reporting.

```
SI 4091(config)# logging log lldp
```

5. Save the configuration.

```
SI 4091(config)# copy running-config startup-config
```

6. Verify the configuration settings

```
SI 4091(config)# show lldp
```

7. View remote device information as needed.

```
SI 4091(config)# show lldp remote-device [<index number>]
```

Chapter 20. Simple Network Management Protocol

Enterprise NOS provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software.

SNMP Version 1

To access the SNMP agent on the SI4091, the read and write community strings on the SNMP manager should be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands on the CLI:

```
SI 4091(config)# snmp-server read-community <1-32 characters>
   -and-
SI 4091(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager should be able to reach the management interface or any one of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch with the following command:

```
SI 4091(config)# snmp-server trap source <trap source IP interface>
SI 4091(config)# snmp-server host <IPv4 address> <trap host community string>
```

SNMP Version 3

SNMP version 3 (SNMPv3) is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMPv3 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators and encryption to protect against threats such as masquerade, modification of information, message stream modification and disclosure.

SNMPv3 allows clients to query the MIBs securely.

Default Configuration

Enterprise NOS has SNMPv3 disabled by default. If a user-created SNMPv3 user is found on the system, SNMPv3 is enabled for backwards compatibility.

Up to 17 SNMP users can be configured on the switch. To modify an SNMP user, enter the following commands:

```
SI 4091(config)# snmp-server user <1-17> name <1-32 characters>
```

Users can be configured to use the authentication/privacy options. The SI4091 support two authentication algorithms: MD5 and SHA, as specified in the following command:

```
SI 4091(config)# snmp-server user <1-17> authentication-protocol {md5|sha}
   authentication-password
   -or-
SI 4091(config)# snmp-server user <1-17> authentication-protocol none
```

User Configuration Example

1. To configure a user with name “admin,” authentication type MD5, and authentication password of “admin,” privacy option DES with privacy password of “admin,” use the following CLI commands.

```
SI 4091(config)# snmp-server user 5 name admin
SI 4091(config)# snmp-server user 5 authentication-protocol md5
     authentication-password
Changing authentication password; validation required:
Enter current admin password:         <admin. password>
Enter new authentication password:    <auth. password>
Re-enter new authentication password: <auth. password>
New authentication password accepted.

SI 4091(config)# snmp-server user 5 privacy-protocol des privacy-password
Changing privacy password; validation required:
Enter current admin password:         <admin. password>
Enter new privacy password:           <privacy password>
Re-enter new privacy password:        <privacy password>
New privacy password accepted.
```

2. Configure a user access group, along with the views the group may access. Use the access table to configure the group’s access level.

```
SI 4091(config)# snmp-server access 5 name admingrp
SI 4091(config)# snmp-server access 5 level authpriv
SI 4091(config)# snmp-server access 5 read-view iso
SI 4091(config)# snmp-server access 5 write-view iso
SI 4091(config)# snmp-server access 5 notify-view iso
```

   Because the read view, write view, and notify view are all set to “iso,” the user type has access to all private and public MIBs.

3. Assign the user to the user group. Use the group table to link the user to a particular access group.

```
SI 4091(config)# snmp-server group 5 user-name admin
SI 4091(config)# snmp-server group 5 group-name admingrp
```

   If you want to allow user access only to certain MIBs, see “View-Based Configurations” on page 261.

View-Based Configurations

- Switch `User` equivalent

  To configure an SNMP user equivalent to the switch "user" login, use the following configuration:

```
SI 4091(config)# snmp-server user 4 name usr(Configure the user)

SI 4091(config)# snmp-server access 3 name usrgrp(Configure access group 3)
SI 4091(config)# snmp-server access 3 read-view usr
SI 4091(config)# snmp-server access 3 write-view usr
SI 4091(config)# snmp-server access 3 notify-view usr

SI 4091(config)# snmp-server group 3 user-name usr(Assign user to access group
3)
SI 4091(config)# snmp-server group 3 group-name usrgrp

SI 4091(config)# snmp-server view 6 name usr(Create views for user)
SI 4091(config)# snmp-server view 6 tree 1.3.6.1.4.1.1872.2.5.1.2

SI 4091(config)# snmp-server view 7 name usr(Agent information)
SI 4091(config)# snmp-server view 7 tree 1.3.6.1.4.1.1872.2.5.1.3

SI 4091(config)# snmp-server view 8 name usr(L2 statistics)
SI 4091(config)# snmp-server view 8 tree 1.3.6.1.4.1.1872.2.5.2.2

SI 4091(config)# snmp-server view 9 name usr(L2 information)
SI 4091(config)# snmp-server view 9 tree 1.3.6.1.4.1.1872.2.5.2.3

SI 4091(config)# snmp-server view 10 name usr(L3 statistics)
SI 4091(config)# snmp-server view 10 tree 1.3.6.1.4.1.1872.2.5.2.3

SI 4091(config)# snmp-server view 11 name usr(L3 information)
SI 4091(config)# snmp-server view 11 tree 1.3.6.1.4.1.1872.2.5.3.3
```

- Switch `Oper` equivalent

```
SI 4091(config)# snmp-server user 5 name usr(Configure the user)

SI 4091(config)# snmp-server access 4 name opergrp(Configure access group 3)
SI 4091(config)# snmp-server access 4 read-view oper
SI 4091(config)# snmp-server access 4 write-view oper
SI 4091(config)# snmp-server access 4 notify-view oper

SI 4091(config)# snmp-server group 4 user-name oper(Assign oper to access group
4)
SI 4091(config)# snmp-server group 4 group-name opergrp

SI 4091(config)# snmp-server view 20 name oper(Create views for oper)
SI 4091(config)# snmp-server view 20 tree 1.3.6.1.4.1.1872.2.5.1.2

SI 4091(config)# snmp-server view 21 name oper(Agent information)
SI 4091(config)# snmp-server view 21 tree 1.3.6.1.4.1.1872.2.5.1.3

SI 4091(config)# snmp-server view 22 name oper(L2 statistics)
SI 4091(config)# snmp-server view 22 tree 1.3.6.1.4.1.1872.2.5.2.2

SI 4091(config)# snmp-server view 23 name oper(L2 information)
SI 4091(config)# snmp-server view 23 tree 1.3.6.1.4.1.1872.2.5.2.3

SI 4091(config)# snmp-server view 24 name oper(L3 statistics)
SI 4091(config)# snmp-server view 24 tree 1.3.6.1.4.1.1872.2.5.2.3

SI 4091(config)# snmp-server view 25 name oper(L3 information)
SI 4091(config)# snmp-server view 25 tree 1.3.6.1.4.1.1872.2.5.3.3
```

Secure Audit Logging

Flex System managers may use the authentication and encryption protocols of SNMPv3 to securely audit the switch. The audit logs record activity and severity for the overall system, user, and application processes. These logs can be used to trace a user's actions, monitor switch alerts, and confirm intrusion detection.

Networking OS uses SNMPv3 authorization to forward the logs securely to the management tool via the chassis management module (CMM). The switch supports both pulling the logs via SNMP 'Get' requests and the pushing of event logs via SNMP traps. Supported management tools are xHMC and other (security and information event management) SIEM tools like Qradar.

Note: Audit logging is enabled by default and cannot be disabled. The audit logs are accessed remotely via SNMPv3 hosts.

Use the following commands to locally manage the logs:

```
SI 4091(config)# show sal reverse      (Display most recent logs first)
SI 4091(config)# clear sal             (Clear audit logs)
```

Configuring SNMP Trap Hosts

SNMPv1 Trap Host

1. Configure a user with no authentication and password

```
SI 4091(config)# snmp-server user 10 name v1trap
```

2. Configure an access group and group table entries for the user. Use the following menu to specify which traps can be received by the user:

```
SI 4091(config)# snmp-server access <user number>
```

In the following example the user will receive the traps sent by the switch.

```
SI 4091(config)# snmp-server access 10(Access group to view SNMPv1 traps)
     name v1trap
     security snmpv1
     notify-view iso
SI 4091(config)# snmp-server group 10  (Assign user to the access group)
     security snmpv1
     user-name v1trap
     group-name v1trap
```

3. Configure an entry in the notify table. .

```
SI 4091(config)# snmp-server notify 10 name v1trap
SI 4091(config)# snmp-server notify 10 tag v1trap
```

4. Specify the IPv4 address and other trap parameters in the `targetAddr` and `targetParam` tables. Use the following menus to specify the user name associated with the targetParam table: :

```
SI 4091(config)# snmp-server target-address 10 name v1trap address
     10.70.70.190
SI 4091(config)# snmp-server target-address 10 parameters-name v1param
SI 4091(config)# snmp-server target-address 10 taglist v1param
SI 4091(config)# snmp-server target-parameters 10 name v1param
SI 4091(config)# snmp-server target-parameters 10 user-name v1only
SI 4091(config)# snmp-server target-parameters 10 message snmpv1
```

Note: Enterprise NOS 8.4 supports only IPv4 addresses for SNMP trap hosts.

5. Use the community table to specify which community string is used in the trap..

```
SI 4091(config)# snmp-server community 10(Define the community string)
     index v1trap
     name public
     user-name v1trap
```

SNMPv2 Trap Host Configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model, use `snmpv2` instead of `snmpv1`. .

```
SI 4091(config)# snmp-server user 10 name v2trap

SI 4091(config)# snmp-server group 10 security snmpv2
SI 4091(config)# snmp-server group 10 user-name v2trap
SI 4091(config)# snmp-server group 10 group-name v2trap
SI 4091(config)# snmp-server access 10 name v2trap
SI 4091(config)# snmp-server access 10 security snmpv2
SI 4091(config)# snmp-server access 10 notify-view iso

SI 4091(config)# snmp-server notify 10 name v2trap
SI 4091(config)# snmp-server notify 10 tag v2trap

SI 4091(config)# snmp-server target-address 10 name v2trap
     address 100.10.2.1
SI 4091(config)# snmp-server target-address 10 taglist v2trap
SI 4091(config)# snmp-server target-address 10 parameters-name v2param
SI 4091(config)# snmp-server target-parameters 10 name v2param
SI 4091(config)# snmp-server target-parameters 10 message snmpv2c
SI 4091(config)# snmp-server target-parameters 10 user-name v2trap
SI 4091(config)# snmp-server target-parameters 10 security snmpv2

SI 4091(config)# snmp-server community 10 index v2trap
SI 4091(config)# snmp-server community 10 user-name v2trap
```

Note: Enterprise NOS 8.4 supports only IPv4 addresses for SNMP trap hosts.

SNMPv3 Trap Host Configuration

To configure a user for SNMPv3 traps, you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

This is configured in the access table using the following commands: :

```
SI 4091(config)# snmp-server access <1-32> level
SI 4091(config)# snmp-server target-parameters <1-16>
```

Configure the user in the user table accordingly.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user `v3trap` with authentication only: :

```
SI 4091(config)# snmp-server user 11 name v3trap
SI 4091(config)# snmp-server user 11 authentication-protocol md5
     authentication-password
Changing authentication password; validation required:
Enter current admin password:         <admin. password>
Enter new authentication password:    <auth. password>
Re-enter new authentication password: <auth. password>
New authentication password accepted.
SI 4091(config)# snmp-server access 11 notify-view iso
SI 4091(config)# snmp-server access 11 level authnopriv
SI 4091(config)# snmp-server group 11 user-name v3trap
SI 4091(config)# snmp-server group 11 tag v3trap
SI 4091(config)# snmp-server notify 11 name v3trap
SI 4091(config)# snmp-server notify 11 tag v3trap
SI 4091(config)# snmp-server target-address 11 name v3trap address
     47.81.25.66
SI 4091(config)# snmp-server target-address 11 taglist v3trap
SI 4091(config)# snmp-server target-address 11 parameters-name v3param
SI 4091(config)# snmp-server target-parameters 11 name v3param
SI 4091(config)# snmp-server target-parameters 11 user-name v3trap
SI 4091(config)# snmp-server target-parameters 11 level authNoPriv
```

Note: Enterprise NOS 8.4 supports only IPv4 addresses for SNMP trap hosts.

SNMP MIBs

The Enterprise NOS SNMP agent supports SNMP version 3. Security is provided through SNMP community strings. The default community strings are "`public`" for `SNMP GET` operation and "`private`" for `SNMP SET` operation. The community string can be modified only through the Command Line Interface (CLI). Detailed SNMP MIBs and trap definitions of the Enterprise NOS SNMP agent are contained in the following Enterprise NOS enterprise MIB document:

```
GbFSIM-10G-L2.mib
```

The Enterprise NOS SNMP agent supports the following standard MIBs:

- `dot1x.mib`
- `ieee8021ab.mib`
- `ieee8023ad.mib`
- `lldpxdcbx.mib`
- `rfc1213.mib`
- `rfc1215.mib`
- `rfc1493.mib`
- `rfc1573.mib`
- `rfc1643.mib`
- `rfc1657.mib`
- `rfc1850.mib`
- `rfc1907.mib`
- `rfc2037.mib`
- `rfc2233.mib`
- `rfc2465.mib`
- `rfc2571.mib`
- `rfc2572.mib`
- `rfc2573.mib`
- `rfc2574.mib`
- `rfc2575.mib`
- `rfc2576.mib`
- `rfc3176.mib`

The Enterprise NOS SNMP agent supports the following generic traps as defined in RFC 1215:

- ColdStart
- WarmStart
- LinkDown
- LinkUp
- AuthenticationFailure

The following are the enterprise SNMP traps supported in Enterprise NOS:

Table 22. *Enterprise NOS-Supported Enterprise SNMP Traps*

Trap Name	Description
altSwDefGwUp	Signifies that the default gateway is alive.
altSwDefGwDown	Signifies that the default gateway is down.
altSwDefGwInService	Signifies that the default gateway is up and in service
altSwDefGwNotInService	Signifies that the default gateway is alive but not in service
altSwLoginFailure	Signifies that someone failed to enter a valid username/password combination.
altSwTempExceedThreshold	Signifies that the switch temperature has exceeded maximum safety limits.
altSwTempReturnThreshold	Signifies that the switch temperature has returned below maximum safety limits.
altSwHotlinksMasterUp	Signifies that the Master interface is active.
altSwHotlinksMasterDn	Signifies that the Master interface is not active.
altSwHotlinksBackupUp	Signifies that the Backup interface is active.
altSwHotlinksBackupDn	Signifies that the Backup interface is not active.
altSwHotlinksNone	Signifies that there are no active interfaces.
altSwValidLogin	Signifies that a user login has occurred.
altSwValidLogout	Signifies that a user logout has occurred.

Switch Images and Configuration Files

This section describes how to use MIB calls to work with switch images and configuration files. You can use a standard SNMP tool to perform the actions, using the MIBs listed in Table 23.

Table 23 lists the MIBS used to perform operations associated with the Switch Image and Configuration files.

Table 23. *MIBs for Switch Image and Configuration Files*

MIB Name	MIB OID
agTransferServer	1.3.6.1.4.1872.2.5.1.1.7.1.0
agTransferImage	1.3.6.1.4.1872.2.5.1.1.7.2.0
agTransferImageFileName	1.3.6.1.4.1872.2.5.1.1.7.3.0
agTransferCfgFileName	1.3.6.1.4.1872.2.5.1.1.7.4.0
agTransferDumpFileName	1.3.6.1.4.1872.2.5.1.1.7.5.0
agTransferAction	1.3.6.1.4.1872.2.5.1.1.7.6.0
agTransferLastActionStatus	1.3.6.1.4.1872.2.5.1.1.7.7.0
agTransferUserName	1.3.6.1.4.1872.2.5.1.1.7.9.0
agTransferPassword	1.3.6.1.4.1.1872.2.5.1.1.7.10.0
agTransferTSDumpFileName	1.3.6.1.4.1.1872.2.5.1.1.7.11.0

The following SNMP actions can be performed using the MIBs listed in Table 23.

- Load a new Switch image (boot or running) from an FTP/SFTP/TFTP server
- Load a previously saved switch configuration from an FTP/SFTP/TFTP server
- Save the switch configuration to an FTP/SFTP/TFTP server
- Save a switch dump to an FTP/SFTP/TFTP server

Loading a New Switch Image

To load a new switch image with the name "`MyNewImage-1.img`" into `image2`, follow these steps. This example shows an FTP/SFTP/TFTP server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/SFTP/TFTP server address where the switch image resides:

   ```
   Set agTransferServer.0 "192.168.10.10"
   ```

2. Set the area where the new image will be loaded:

   ```
   Set agTransferImage.0 "image2"
   ```

3. Set the name of the image:

   ```
   Set agTransferImageFileName.0 "MyNewImage-1.img"
   ```

4. If you are using an FTP/SFTP server, enter a username:

   ```
   Set agTransferUserName.0 "MyName"
   ```

5. If you are using an FTP/SFTP server, enter a password:

   ```
   Set agTransferPassword.0 "MyPassword"
   ```

6. Initiate the transfer. To transfer a switch image, enter 2 (gtimg):

   ```
   Set agTransferAction.0 "2"
   ```

Loading a Saved Switch Configuration

To load a saved switch configuration with the name "MyRunningConfig.cfg" into the switch, follow these steps. This example shows a file server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/SFTP/TFTP server address where the switch Configuration File resides:

   ```
   Set agTransferServer.0 "192.168.10.10"
   ```

2. Set the name of the configuration file:

   ```
   Set agTransferCfgFileName.0 "MyRunningConfig.cfg"
   ```

3. If you are using an FTP/SFTP server, enter a username:

   ```
   Set agTransferUserName.0 "MyName"
   ```

4. If you are using an FTP/SFTP server, enter a password:

   ```
   Set agTransferPassword.0 "MyPassword"
   ```

5. Initiate the transfer. To restore a running configuration, enter 3:

   ```
   Set agTransferAction.0 "3"
   ```

Saving the Switch Configuration

To save the switch configuration to a FTP/SFTP/TFTP server follow these steps. This example shows a file server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/SFTP/TFTP server address where the configuration file is saved:

   ```
   Set agTransferServer.0 "192.168.10.10"
   ```

2. Set the name of the configuration file:

   ```
   Set agTransferCfgFileName.0 "MyRunningConfig.cfg"
   ```

3. If you are using an FTP/SFTP server, enter a username:

   ```
   Set agTransferUserName.0 "MyName"
   ```

4. If you are using an FTP/SFTP server, enter a password:

   ```
   Set agTransferPassword.0 "MyPassword"
   ```

5. Initiate the transfer. To save a running configuration file, enter 4:

   ```
   Set agTransferAction.0 "4"
   ```

Saving a Switch Dump

To save a switch dump to a FTP/SFTP/TFTP server, follow these steps. This example shows an file server at 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/SFTP/TFTP server address where the configuration will be saved:

   ```
   Set agTransferServer.0 "192.168.10.10"
   ```

2. Set the name of dump file:

   ```
   Set agTransferDumpFileName.0 "MyDumpFile.dmp"
   ```

3. If you are using an FTP/SFTP server, enter a username:

   ```
   Set agTransferUserName.0 "MyName"
   ```

4. If you are using an FTP/SFTP server, enter a password:

   ```
   Set agTransferPassword.0 "MyPassword"
   ```

5. Initiate the transfer. To save a dump file, enter 5:

   ```
   Set agTransferAction.0 "5"
   ```

Chapter 21. Secure Input/Output Module

The Secure Input/Output Module (SIOM) enables you to determine which protocols can be enabled. The SIOM only allows secured traffic and secured authentication management.

The following topics are discussed in this chapter:

SIOM Overview

In networking solutions, a new approach about adopting a security level on Input/Output modules has been developed. This security level encompasses secured authentication management and only allows secure traffic and protocols.

IOMs can be classified into two security categories:

- Legacy Input/Output Modules (LIOMs)

  LIOMs are not capable of provisioning any security policy setting. All IOMs developed before the SIOM feature was introduced are of type LIOM.

- Secure Input/Output Modules (SIOMs)

  SIOMs have security characteristics that allow them to integrate the network assigned security policy.

For IOM to be in SIOM mode, both the IOM and the CMM (Chassis Management Module) containing it must be running SIOM-capable software. In all other cases, the IOM operates in LIOM mode.

When the IOM is in SIOM mode, the security characteristics configured on the CMM are sent to the IOM. These characteristics can be divided into the following categories:

- Policy setting
- User Account Management
- Secure LDAP (LDAPS) authentication

SIOM is disabled by default. To enable SIOM on the switch, in Global Configuration mode, enter:

```
SI 4091(config)# boot siom enable
```

To disable SIOM, enter:

```
SI 4091(config)# no boot siom enable
```

Note: You must reboot the switch for SIOM settings to take effect.

Switch Access in SIOM Mode

After the embedded switch is provisioned by the CMM in the SIOM mode, the switch will automatically update its LDAP settings (`startTLS`, `LDAPS` or `LDAP`) to the ones configured on the CMM. When no external LDAP server is configured on the CMM, CMM itself will serve as the local LDAP server. The LDAP client configured on the CMM is pushed onto the switch and the LDAP credentials used to access the CMM can also be used to access the switch.

To access the switch, you may now use one of the following methods:

- The CMM credentials
- Other user credentials which depend on the SIOM security policy setting, as follows:
  - In legacy mode, if RADIUS or TACACS+ is enabled, they will replace LDAP as authentication method. If LDAP local access mode is enabled, you can still use local authentication by using `noldap` as username.
  - In secure mode, you may use the provisioned LDAP credentials.

Note: Once the switch is provisioned by the CMM in the SIOM mode, it cannot be accessed using the switch local user accounts.

SIOM in Stacking Mode

In stacking mode, configuring SIOM is only supported on the Master switch. Hence, the command:

```
SI 4091(config)# [no] boot siom enable
```

is only supported on the Master switch. On stack member switches, SIOM is configured by the Master switch, and the member switches automatically inherit the Master switch SIOM setting. When upgrading to SIOM-capable software:

- The Master, Backup, and member switches need to be rebooted for SIOM to take effect.
- When SIOM is enabled on the Master, it is applied on all stack members automatically.
- If a new switch with a different boot SIOM configuration is attached to the stack, the switch will inherit the boot SIOM configuration from the Master and will automatically reboot.
- When two stacks are joined, the selected Master for the two stacks will push its own boot SIOM configuration, and the added members will automatically reboot.
- There will be no changes in the SIOM policy on members if the stack is split.

Note: Lenovo recommends using staggered upgrade. In this case, the upgrade will take more time, depending on how large the stack setup is, but the traffic loss will be minimal.

Overall, in a stacking mode, it takes longer (about 9 minutes) to reload SIOM-enabled software than a non-SIOM enabled software. If staggered-upgrade procedure is used, this duration increases according to the number of switches in the stack.

The process of upgrading from non-SIOM enabled software to SIOM-enabled software takes about 15 minutes. If a staggered-upgrade procedure is used, this duration increases according to the number of switches in the stack.

If the Master switch gets rebooted, the Backup switch becomes the Master (operation called Master failover) and it will be SIOM provisioned. If the SIOM provisioning occurs for the first time on this switch, it will also reboot for the LIOM to SIOM transition. The new Master will reboot only after the Backup switch (original Master) rejoins the stack.

SIOM Feature Considerations

SIOM has two aspects which must be accomplished on the switch:

- The provisioning process which supplies all the necessary settings for the secure Ethernet connection for management and for the secure protocols enabled on the switch.
- The protocols enabled during the functioning of a switch in SIOM.

Switch boots up with all operational (data) ports disabled. Although the management ports are enabled, they can't be used by `admin` to set up the switch until the configuration is applied.

Internal management port is used by the CMM during the provisioning to exchange information with IOM.

At the end of provisioning, when SIOM is enabled, the rest of the operational ports come up and the switch will be fully functional.

When in SIOM mode, the PKI system of switch cannot be controlled. The user cannot import his own certificate. All certificates are provisioned by CMM.

Creating a Policy Setting

The policy setting can be either secure (IOM is in secure mode) or legacy (IOM is in legacy mode). In secure mode, only communication protocols that are deemed secure can be used; most protocols that are not deemed secure are disabled. In legacy mode setting, all protocols are allowed (LIOM behavior).

To display the current policy setting, enter:

```
SI 4091(config)# show boot security-policy
```

Note: You must reboot the IOM for a new policy setting to be applied.

Protocols Affected by the Policy Setting

This section explains which protocols can and cannot operate in secure mode on the SI4091 System Interconnect Module.

Insecure Protocols

When you are in Secure Mode, the following protocols are deemed "insecure" and are disabled:

- HTTP
- LDAP Client
- SNMPv1
- SNMPv2
- Telnet (server and client)
- FTP (server and client)
- Radius (client
- TFTP Server

Except for the TFTP server, these protocols cannot be enabled when the switch is operating in Secure Mode because the commands to enable or disable them are no longer enabled.

The following protocols, although deemed "insecure," are enabled by default and can be disabled.

- DHCP client
- SysLog

Note: Service Location Protocol (SLP) Discovery is also deemed "insecure" but is unaffected by Secure Mode. SLP has the same default settings as in Legacy Mode. If you can enable or disable SLP in Legacy Mode, you can enable or disable it the same way in Secure Mode.

The following supported protocols are not enabled by default but can always be enabled in Secure Mode.

- DNS Resolution
- TFTP client (for signed items only, such as switch images)

Secure Protocols

The following protocols are deemed “secure” and are enabled by default in Secure Mode:

- SCP Server
- SNMPv3 Client
- SFTP Client
- SSHv2 Server
- SSHv2 Client
- HTTPS Server
- TACACS+ (client)

You can disable these protocols.

The following protocols are deemed “secure” and cannot be disabled in any mode:

- NTP Client v4
- LDAPS Client

The following protocols are also deemed “secure” on the SI4091 and can be enabled.

- IKE
- IPSec

The default state for these protocols in Secure Mode, whether enabled or disabled, is the same as in Legacy Mode.

The following protocols are deemed “secure” but are not currently supported by the SI4091:

- EAPoL
- SCP
- S/MIME
- SNMPv3 Manager
- TCP command secure mode (Port 6091)

Insecure Protocols Unaffected by SIOM

The following protocols are deemed “insecure” but can be enabled in all Security Policy Modes:

- Ping
- Ping IPv6
- Traceroute
- Traceroute IPv6
- TFTP IPv6
- SNMPv3 IPv6

- bootp

Notes:

- Telnet IPv6 and TFTP IPv6 are disabled in Secure Mode.
- TFTP IPv6 is allowed in Secure Mode for signed image transfers only.

Managing User Accounts

SNMPv3 user accounts with customized attributes can be created on the CMM and pushed to the IOM. For each SNMPv3 user account created on the CMM, the IOM creates a local SNMPv3 user account. The SNMPv3 user database then creates new user-per-profile user lists. It then uses this database to authenticate users.

Note: SNMPv3 does not support LDAP user management, so the CMM must provision SNMPv3 user accounts to the IOM.

Using Centralized SNMPv3 Management with SIOM

There is a setting on the CMM to indicate whether the SNMPv3 centralized user management is enabled; this is called the *Centralized Flag*.

When the IOM runs in SIOM mode and the Centralized Flag is enabled, SNMPv3 will enable Node Accounts and will disable Local Accounts. When the IOM runs in LIOM mode *or* the Centralized Flag is disabled, SNMPv3 will use Local Accounts and disable Node Accounts. Node Accounts represent accounts configured on the CMM, while Local Accounts are accounts configured on the IOM.

Since there is no case where both the Node Account and Local Account are enabled, the username of a Node Account can be duplicated with a Local Account username.

Implementing SNMPv3 with SIOM

The following commands are available for implementing SNMPv3 with SIOM:

- `access snmp read-only`
- `access snmp read-write`
- `snmp-server access`
- `snmp-server community`
- `snmp-server group`
- `snmp-server host`
- `snmp-server notify`
- `snmp-server read-community`
- `snmp-server read-community-additional`
- `snmp-server target-address`
- `snmp-server target-parameters`
- `snmp-server user`
- `snmp-server version`
- `snmp-server view`
- `snmp-server write-community`
- `snmp-server write-community-additional`
- `show snmp-server v3`

For more information about these commands, see the *Lenovo ISCLI—Industry Standard CLI Command Reference for the Lenovo Flex System SI4091 10Gb System Interconnect Module*.

Implementing Secure LDAP (LDAPS)

Lightweight Directory Access Protocol (LDAP) is a protocol for accessing distributed directory information services over a network. Enterprise NOS uses LDAP for authentication and authorization. With an LDAP client enabled, the switch will authenticate a user and determine the user's privilege level by checking with one or more directory servers instead of a local database of users. This prevents customers from having to configure local user accounts on multiple switches; they can maintain a centralized directory instead.

As part of SIOM, you can implement Secure Lightweight Directory Access Protocol (LDAPS) in addition to standard LDAP.

Enabling LDAPS

When the IOM is in SIOM mode, all LDAP configurations are made from the CMM and pushed to the IOM. When the IOM is in LIOM mode, the CLI can be used to configure LDAP settings. LDAPS is disabled by default. To enable LDAPS:

1. Turn LDAP authentication on

```
SI 4091(config)# ldap-server enable
```

2. Enable LDAP Enhanced Mode:

```
SI 4091(config)# ldap-server mode enhanced
```

This changes the `ldap-server` subcommands to support LDAPS.

3. Configure the IPv4 addresses of each LDAP server.

```
SI 4091(config)# ldap-server host {1-4} <IP address or hostname>
```

4. You may change the default TCP port number used to listen to LDAPS (optional).

The well-known port for LDAP is 636.

```
SI 4091(config)# ldap-server port <1-65000>
```

5. Configure the Security Mode:

```
SI 4091(config)# ldap-server security {clear|ldaps|mutual|starttls}
```

where:

Parameter	Description
clear	Cleartext Mode (no security)
ldaps	LDAPS Mode
mutual	Mutual authentication in Transport Layer Security (TLS)
starttls	Secure LDAP via StartTLS without cleartext fallback

6. Configure the distinguished name (DN) and password (optional).

```
SI 4091(config)# ldap-server binddn dn "<distinguished name> "
SI 4091(config)# ldap-server binddn key "<password> "
```

If this is not configured, the switch will use user-provided login credentials to bind. A DN will then be constructed from the user's login credentials and then used in the initial BIND attempt.

7. Configure the root DN:

```
SI 4091(config)# ldap-server basedn <root DN name>
```

8. Configure the user search attribute (optional):

```
SI 4091(config)# ldap-server attribute username <search attribute>
```

If no user search attribute is specified, the default is `uid`.

9. Configure the group search attribute (optional):

```
SI 4091(config)# ldap-server attribute group <search attribute>
```

If no group search attribute is specified, the default is `memberOf`.

10. Configure the login permissions attribute:

```
SI 4091(config)# ldap-server attribute login-permission <attribute>
```

Note: If no login permissions attribute is configured, LDAP client will not function.

11. Configure the group filter attribute (optional):

```
SI 4091(config)# ldap-server group-filter <filter attributes separated by comma>
```

Note: The group filter string must contain no whitespace.

If no group filter attribute is configured, no groups will be filtered and all groups will be considered in any search.

12. Enable DNS server verification:

```
SI 4091(config)# ldap-server srv
```

Disabling LDAPS

To disable LDAPS, enter:

```
SI 4091(config)# ldap-server security clear
SI 4091(config)# ldap-server mode legacy
```

For information about using LDAP in Legacy Mode, see "LDAP Authentication and Authorization" on page 99.

Syslogs and LDAPS

Syslogs are displayed for the following error conditions:

- `Password change required on first login`
- `Password expired`
- `Username or password invalid`
- `Account temporarily locked`
- `Unknown/no reason given`

SIOM Dependencies

The following points are relevant to SIOM:

- The CMM has a Certificate Authority (CA) capable of signing the certificates involved for authenticating the IOM in the SSL and TLS processes and protocols.
- The correctness of the configuration depends upon the settings on the CMM. This is especially important for NTP and LDAP, which ensure switch operability. For example, if the LDAP client is configured incorrectly, the switch cannot be managed.
- The Enhanced Configuration and Management (EHCM) module configures the NTP client. Therefore, the NTP client is dependent upon the ECHM module being enabled and functional.
- Some protocols cannot be changed from enabled to disabled without restarting the switch. The IOM may reboot when switching between the SIOM and LIOM.

Part 8: Appendices

Appendix A. Getting Help and Technical Assistance

If you need help, service, or technical assistance or just want more information about Lenovo products, you will find a wide variety of sources available from Lenovo to assist you. This appendix will help you obtain additional information about Lenovo and Lenovo products, and determine what to do if you experience a problem with your Lenovo system or optional device.

Note: This section includes references to IBM web sites and information in regard to obtaining service. IBM is Lenovo's preferred service provider for the System x, Flex System, and NeXtScale System products.

You can solve many problems without outside assistance by following the troubleshooting procedures that Lenovo provides in the online help or in the Lenovo product documentation. The Lenovo product documentation also describes the diagnostic tests that you can perform. The documentation for most systems, operating systems, and programs contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

Before you call, make sure that you have taken these steps to try to resolve the problem yourself.

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Check for updated software, firmware and operating-system device drivers for your Lenovo product. The Lenovo Warranty terms and conditions state that you, the owner of the Lenovo product, are responsible for maintaining and updating all software or firmware for the product (unless it is covered by an additional maintenance contract). Your service technician will request that you upgrade your software or firmware if the problem has a documented solution within a software or firmware upgrade.
- If you have installed new hardware or software in your environment, check the IBM ServerProven website to make sure that the hardware and software is supported by your product.
- Go to the IBM Support portal to check for information to help you solve the problem, and for other support options.

- If you believe that you require warranty service for your Lenovo product, gather the following information to provide to the service technician. This data will help the service technician quickly provide a solution to your problem and ensure that you receive the level of service for which you might have contracted. The service technicians will be able to assist you more efficiently if you prepare before you call.
  - Hardware and Software Maintenance agreement contract numbers, if applicable
  - Machine type number (Lenovo 4-digit machine identifier), if applicable
  - Model number
  - Serial number
  - Current system UEFI and firmware levels
  - Other pertinent information such as error messages and logs

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.

Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties.

Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

Lenovo, the Lenovo logo, Flex System, System x, NeXtScale System, and X-Architecture are trademarks of Lenovo in the United States, other countries, or both.

Intel and Intel Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Internet Explorer, Microsoft, and Windows are trademarks of the Microsoft group of companies.

Linux is a registered trademark of Linus Torvalds.

Other company, product, or service names may be trademarks or service marks of others.

Important Notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1 024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard-disk-drive bays with the largest currently supported drives that are available from Lenovo.

Maximum memory might require replacement of the standard memory with an optional memory module.

Each solid-state memory cell has an intrinsic, finite number of write cycles that the cell can incur. Therefore, a solid-state device has a maximum number of write cycles that it can be subjected to, expressed as total bytes written (TBW). A device that has exceeded this limit might fail to respond to system-generated commands or might be incapable of being written to. Lenovo is not responsible for replacement of a device that has exceeded its maximum guaranteed number of program/erase cycles, as documented in the Official Published Specifications for the device.

Lenovo makes no representations or warranties with respect to non-Lenovo products. Support (if any) for the non-Lenovo products is provided by the third party, not Lenovo.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

Index

Symbols

Numerics

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W