

Lenovo Flex System Interconnect Fabric

Solution Guide

for Networking OS 8.3

LenovoTM

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.

First Edition (November 2015)

© Copyright Lenovo 2015

Portions © Copyright IBM Corporation 2014.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Lenovo and the Lenovo logo are trademarks of Lenovo in the United States, other countries, or both.

Contents

Preface5
Who Should Use This Guide	6
Typographic Conventions	7
Additional References	8
Chapter 1. Introduction9
SI Fabric Architecture10
Point of Delivery11
Hyper Distributed Fabric Protocol.11
SI Fabric Components12
SI4093/EN4093R Switches12
G8264CS Switches.12
SI Fabric Features.13
Layer-2 Switching Capabilities13
Network Simplification13
Management Simplification13
Storage Integration14
Scalable Point of Delivery14
Limitations15
Chapter 2. Initial Configuration17
Hardware Planning18
Initial System Configuration Overview19
Using the ISCLI.19
Before Configuring the SI Fabric19
Converting Standalone Switches to SI Fabric Members Operation20
Specifying SI Fabric Settings on the G8264CS21
Connecting and Attaching the SI Fabric Devices23
Configuring and Binding the SI4093/EN4093R Switches25
Reverting SI Fabric Members to Standalone Operation.27
Configuring the SI Fabric Management IP Addresses and Gateway28
Floating IP Address30
SNMP, Telnet, and SSH Access31
SNMP and Switch Center Access31
Telnet Access.32
SSH Access.32
Displaying SI Fabric Configuration33
Chapter 3. Managing the Flex SI Fabric Solution35
SI4093 and EN4093R License Activation Keys36
Obtaining Activation Keys36
Installing Activation Keys37
Transferring Activation Keys37

Managing Systems through Master and Backup G8264CS Switches	38
Deploying Staggered Image Upgrades for Master G8264CS Switch	38
Deploying Staggered Image Reboot for Master G8264CS Switch.	38
Flexible Port Mapping	39
Microburst Detection	39
Connecting Uplink to the Data Center Network	40
Defining Server Ports	41
Connecting Downlink Compute Nodes	42
Connecting to Storage Resources	44
FIP Snooping Bridge	44
NPV Gateway Mode	45
Full-Fabric Mode	46
Chapter 4. Flex SI Fabric -Specific ISCLI Commands	49
Information Commands	49
Configuration Commands	50
IP Interface Configuration Commands	51
Software Key Commands	52
Boot Options.	53
Appendix A. Getting Help and Technical Assistance	55

Preface

The *Lenovo Flex SI Fabric Solution Guide* describes how to configure and use the Lenovo Networking OS 8.3 software on the Lenovo System Networking RackSwitch G8264CS (referred to as G8264CS throughout this document), the Lenovo Flex System Fabric SI4093 System Interconnect Module (referred to as SI4093) and the Lenovo Flex System Fabric EN4093R 10Gb Scalable Switch (referred to as EN4093R) to achieve the Flex System Interconnect Fabric (referred to as the SI Fabric).

For documentation about the physical installation of the SI4093, EN4093R or G8264CS, see the respective product hardware documentation.

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.

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 <code>readme.txt</code> file.
	It also depicts on-screen computer output and prompts.	<code>host#</code>
ABC123	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	<code>host# sys</code>
<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: <code>host# telnet <IP address></code>
	This also shows book titles, special terms, or words to be emphasized.	Read your <i>User's Guide</i> thoroughly.
[]	Command items shown inside squared brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	<code>host# ls [-a]</code>
{ }	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.	<code>host# set {left right}</code>
AaBbCc123	This block type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the Save button.

Additional References

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

- *Lenovo RackSwitch G8264CS Application Guide for Networking OS 8.3*
- *Lenovo RackSwitch G8264CS ISCLI Command Reference for Networking OS 8.3*
- *Lenovo Flex System Fabric SI4093 System Interconnect Module Application Guide for Networking OS 8.3*
- *Lenovo Flex System Fabric SI4093 System Interconnect Module ISCLI Command Reference for Networking OS 8.3*
- *Lenovo Flex System Fabric EN4093R 10Gb Scalable Switch Application Guide for Networking OS 8.3*
- *Lenovo Flex System Fabric EN4093R 10Gb Scalable Switch ISCLI Command Reference for Networking OS 8.3*

Chapter 1. Introduction

This *Solution Guide* introduces the Lenovo Flex System Interconnect Fabric (referred to as the SI Fabric throughout this document). The SI Fabric provides a single point of network access for data center Ethernet and Fibre Channel network. It supports Flex System compute nodes and internal and external storage nodes for computing and storage resources capabilities in a Flex System chassis.

SI Fabric Architecture

The SI Fabric is managed through a single entity, the user-designated Master switch, and passed on to the internal components. The SI Fabric integrates the following components:

- Two G8264CS switches
- Up to 9 Flex System chassis, each with two Flex modules (SI4093/EN4093R) in paired bays (Bay 1 and Bay 2, or alternately Bay 3 and Bay 4), for a total of up to 18 SI4093/EN4093R modules.

Note: If the servers have an extra net card, the chassis supports up to 4 switches.

- Up to 126 half-width Flex System compute nodes (up to 14 per chassis) with 10Gb Converged Network Adapters (CNAs)
- One or more internal Flex System V7000 storage node and external Storwize V7000 storage system, as a highly integrated Point Of Delivery (POD) solution for data centers

Figure 1. Flex System Interconnect Fabric Components

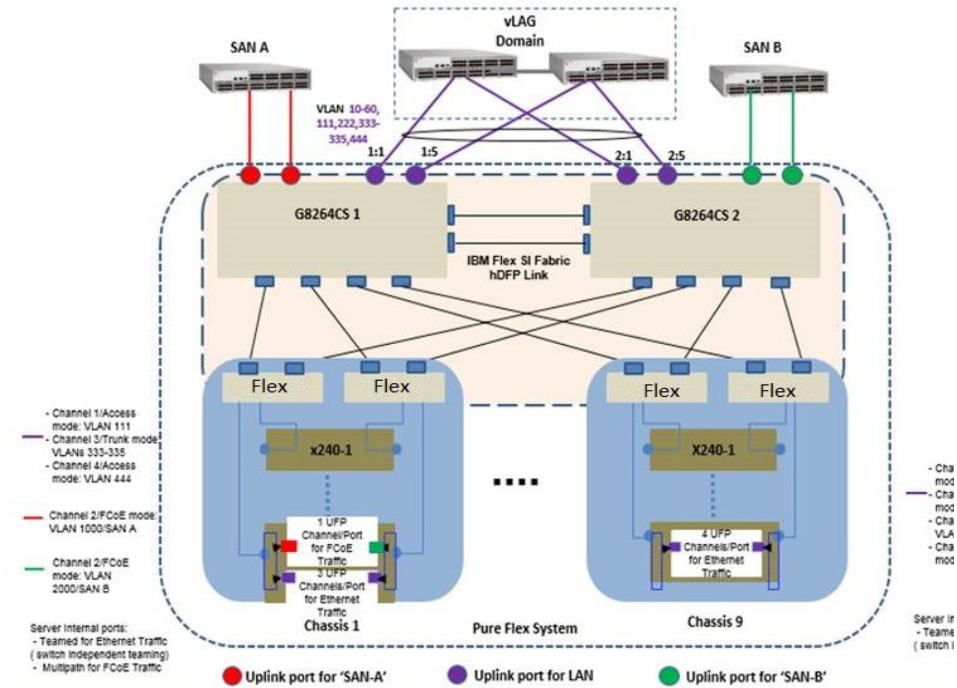
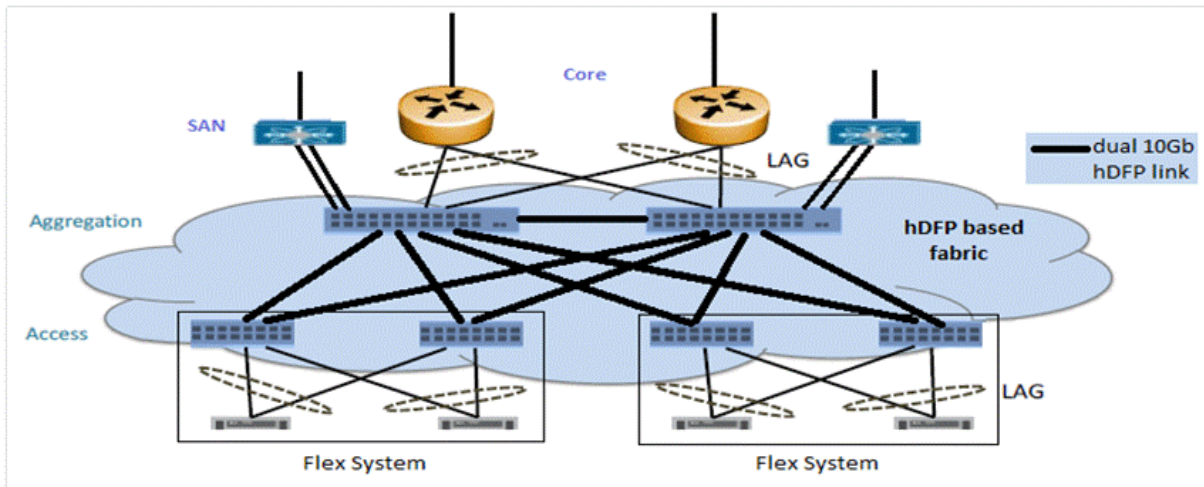


Figure 2 on page 11 illustrates the SI Fabric connections to storage and computing resources.

Figure 2. Flex System Interconnect Fabric Environment



Point of Delivery

A Point Of Delivery (POD) is a module of network, compute, storage, and application components that work together to deliver networking services. In service provider infrastructure, the POD supports cloud computing services to sustain scalability as usage grows.

The SI Fabric integrates the entire POD solution into a seamless network fabric for server and storage under single IP management and attaches to upstream data center network as a loop-free layer-2 stub network fabric with single Ethernet uplink connection or trunk group to each layer-2 network.

The SI Fabric POD solution requires network provisioning only for uplink connections to a data center network, downlink connections to server nodes, and connections to storage nodes.

Hyper Distributed Fabric Protocol

The SI Fabric utilizes the Lenovo proprietary hyper distributed fabric protocol (hDFP) to provide a loop-free network fabric system for data and storage traffic. It integrates compute, network, and storage components as a single IP management switch system. All G8264CS switches and SI4093/EN093R switches are configured to run in SI Fabric mode.

The hDFP aggregates all network switching modules into a single logical switch, the SI Fabric, with G8264CS switches as master and backup nodes at the aggregation layer and SI4093/EN4093R switches as member nodes at the access layer.

Configuration and management of the networking component is done via the management IP interface at master node of the logical switch. If the master node fails, the backup G8264CS switch will take over as the new master node of the logical switch and the old master node will turn into new backup node once recovered.

SI Fabric Components

The SI Fabric unifies SI4093, EN4093R, G8264CS, and storage hardware components.

SI4093/EN4093R Switches

The Lenovo Flex System Fabric SI4093 System Interconnect Module (SI4093) and Lenovo Flex System Fabric EN4093R 10Gb Scalable Switch (EN4093R) provide a transparent network interface between server elements within the chassis and the upstream data and storage networking infrastructure. The standard SI4093 and EN4093R provide 14 internal 10GbE ports as downlinks for server connectivity and 10 external 10GbE Ethernet fabric ports for connection with G8264CS switches.

The SI Fabric can also support expansion of additional 28 internal and 12 external 10Gbps ports, with optional Feature on Demand (FoD) license keys. Within the SI Fabric, SI4093 and EN4093R provide convergence of traditional Ethernet traffic with Fibre Channel over Ethernet (FCoE) storage session traffic and facilitates low-latency, media-speed server-to-server traffic within the chassis, or server-to-uplink traffic based on the hDFP protocol.

G8264CS Switches

The Lenovo System Networking RackSwitch G8264CS converged switch contains 36 10GbE small form-factor pluggable plus (SFP+) ports and four 40GbE quad small-form-factor pluggable plus (QSFP+) ports, and twelve Omni Ports. Omni Port technology allows the same port to be used as a 10G Ethernet port or 4/8G Fibre Channel port.

SI Fabric Features

Layer-2 Switching Capabilities

The SI Fabric provides the following layer-2 switching capabilities:

- Converged Enhanced Ethernet (CEE)
- 10GbE Ethernet switching
- Fibre Channel over Ethernet (FCoE) and FC/FCoE capabilities
 - N Port Virtualization (NPV) Gateway
 - Full Fabric FC/FCoE switch
- VLANs
- Double Tagging (tagpvid)
- Distributed static/LACP portchannels
- Distributed MAC learning
- Teaming

Network Simplification

The SI Fabric provides the following network simplification features:

- Seamless network fabric for server and storage to connect with data center
- Offers a loop-free network fabric without Spanning Tree Protocol (STP) complexity.
- Minimizes network latency by local L2 switching at every interconnect component, and minimizes the loss of data during network failover within the fabric.
- Converges Ethernet for lossless storage traffic.
- Adds new components to an already running fabric with no traffic loss or disturbance in the existing traffic flow.
- Integrates Fibre Channel Forwarder (FCF) to provide end-to-end Fibre Channel over Ethernet (FCoE) storage functionality within the POD without the need of expensive Fibre Channel switch.
- Supports single-fabric mode topology, dual-fabric mode topology, and full-fabric topology.
- Offers different levels of redundancy inside the fabric: static/LACP distributed portchannels, teaming, dual interconnection redundancy, hot links, master/backup switch.

Management Simplification

The SI Fabric provides the following management simplification features:

- High availability with master and backup nodes
- Staggered upgrade enables upgrading the fabric with no service or connection downtime.
- Minimizes managed network elements with single point of management at the master node of the entire fabric.

- Establishes clear administrative boundaries in data center by pushing traditional networking configuration outside of the POD.
- Integrates physical and virtual infrastructure management for compute, network, and storage elements.

Storage Integration

The SI Fabric simplifies integration of storage and storage virtualization with IBM Storwize V7000 storage node and access to external SAN storage infrastructure, such as IBM Storwize V7000, V5000 or V3700.

Scalable Point of Delivery

The SI Fabric provides the following scalability features:

- The SI Fabric enables the size of the POD to grow without the additions of management complexity.
- The SI Fabric adds additional chassis resources up to the maximum configuration under single IP management of the POD.
- Expert Integrated Systems
- The SI Fabric simplifies the I/O connectivity in Foundation offerings and improves the time to value of an integrated system.
- The SI Fabric can be leveraged in Pure Application and Pure Data configurations improving the ability to roll in racks of tested solutions into existing data centers.

Limitations

Following is a list of Flex System Interconnect Fabric limitations for the current release:

- The SI Fabric requires Lenovo Networking OS version 8.2 for all G8264CS and SI4093 standalone switches and Lenovo Networking OS version 8.3 for all EN4093 standalone switches.
- Lenovo does not support mixing 10Gb and 40Gb SI Fabric port speeds in SI Fabric logical switches. Lenovo recommends that you reserve additional SI Fabric ports on both G8264CS switches for future expansion.
- All external ports on SI4093/EN4093R devices are automatically configured as Fabric Ports, requiring no additional configuration. Omni Ports on the RackSwitch G8264CS cannot be configured as Fabric Ports.
- To change a port from Ethernet to Fabric mode on the aggregation switches (TORs), the switch needs to be rebooted. If a new chassis is added to the solution and fabric ports need to be reconfigured on both TORs, the user can still operate master/backup failover and reboot the master and backup one by one in order to keep the full access to its servers/storage during the process.

Chapter 2. Initial Configuration

This chapter explains how to set up the software components of the Flex System Interconnect Fabric (SI Fabric).

Hardware Planning

The minimum configuration requires two G8264CS switches, one Flex System chassis populated with two SI4093 or EN4093R switches and two half-width Flex System compute nodes with 10Gb CNAs at each chassis.

The SI Fabric supports two G8264CS switches, up to nine Flex System chassis populated with two SI4093 or EN4093R switches at each chassis, up to 126 half-width Flex System compute nodes and external storage systems.

The SI Fabric acts as a networking provision simplifier. The solution is designed to behave most like an end-host-mode or transparent mode switch, building an active redundant data traffic path between the upper networks and the servers and from one server to another server in the same sets of VLANs.

- The base SI4093 and EN4093R switches provide 14 internal 10GbE ports as downlinks for server connectivity and ten external 10GbE Ethernet fabric ports for connection with G8264CS switches in the SI Fabric.
- An additional 28 internal and 12 external 10Gbps ports are available for expansion with optional Feature on Demand license keys. This applies to each SI4093 or EN4093R module.
- The SI4093 and EN4093R switches in the SI Fabric provide the convergence of traditional Ethernet traffic with Fibre Channel over Ethernet (FCoE) storage.

Initial System Configuration Overview

To deploy the Flex System Interconnect Fabric from a standalone topology:

1. Specify SI Fabric ports and SI Fabric-related settings at G8264CS switches and reboot again to take effect (see [“Specifying SI Fabric Settings on the G8264CS” on page 21](#)).
2. Connect minimum two 10Gb Ethernet cables as SI Fabric links between the two G8264CS switches. Also connect two 10Gb Ethernet cables from each SI4093 or EN4093R switch to each G8264CS switch. See [“Connecting and Attaching the SI Fabric Devices” on page 23](#).
Note: Using more links to connect switches will increase the traffic bandwidth inside the fabric.
3. Once the master G8264CS detects all attached SI4093 or EN4093R switches, bind them to the SI Fabric (see [“Configuring and Binding the SI4093/EN4093R Switches” on page 25](#)).
4. Configure the management IP address and gateway information at the master and backup G8264CS switches for future data and storage VLAN and port configurations (see [“Configuring the SI Fabric Management IP Addresses and Gateway” on page 28](#)).
5. Enable additional management access (see [“SNMP, Telnet, and SSH Access” on page 31](#)).
6. If required, install FoD license keys on SI4093 or EN4093R switches (see [“SI4093 and EN4093R License Activation Keys” on page 36](#)).

Using the ISCLI

Use the ISCLI via switch serial console to configure the SI Fabric. After the SI Fabric is initialized, you can use ISCLI over Telnet, SSH or SNMP for additional configuration and management tasks.

The ISCLI provides a direct method for collecting switch information and performing switch upgrade and configuration. Using a basic terminal, the ISCLI allows you to view information and statistics about the switch, and to perform any necessary configuration.

Before Configuring the SI Fabric

Before you begin configuring the SI Fabric, you must have this information:

- IP addresses for your routers, servers, and attached network devices
- VLAN information (names and numbers)
- Physical layout, storage, and data connections for the SI Fabric

Converting Standalone Switches to SI Fabric Members Operation

You can add to the SI Fabric any standalone G8264CS, SI4093 and EN4093 switches.

Note: By converting to SI fabric, the switch will automatically lose all its previous configuration.

Use the following commands to load a SI Fabric image to a standalone switch:

```
Router> enable
Router# copy {tftp|ftp|sftp} {image1|image2|boot-image} address <IP address>
filename <SI Fabric filename> {data-port|mgt-port}
```

Note: For EN4093R switches running in stacking mode, make sure to disable the stack mode before uploading the SI Fabric image.

Specifying SI Fabric Settings on the G8264CS

After you have rebooted all G8264CS and SI4093 or EN4093R switches with installed SI Fabric software images, they will load factory-default switch settings and default SI Fabric settings.

Perform the following configuration procedure on each G8264CS switch.

Note: Lenovo recommends that you perform side-by-side logical switch setting review and re-configuration at both G8264CS switches.

1. Connect to the G8264CS switch serial console.
2. Assign SI Fabric ports (those that will be connected to the other switches in the SI Fabric):

```
RS G8264CS> enable
RS G8264CS# boot fabric fabric-trunk 1:17-1:38
A Reboot is required for the new settings to take effect.
```

Note: The required reboot may be performed in a later step.

3. Assign a unique domain ID to the SI Fabric:

```
RS G8264CS# boot fabric si-fabric-domain 82
Current fabric domain: 1
New fabric domain: 82
A reboot is required for the new settings to take effect
```

Assign the same SI Fabric domain ID (between 1 and 100) on both G8264CS switches. This will permit them to join the same logical switch, becoming master and backup nodes. The default domain ID is 1.

See [“Boot Options” on page 53](#) for more information and related commands.

4. Verify that one of the G8264CS switches is set as the master switch:

```
RS G8264CS# show boot fabric
Current fabric settings:
Fabric           : ON
Switch Mode      : Master
Fabric Trunk Ports: 1:17-1:35
Fabric VLAN      : 4090
Fabric domain: 82

Saved fabric settings:
Fabric           : ON
Switch Mode      : Master
Fabric Trunk Ports: 1:17-1:35
Fabric VLAN      : 4090
Fabric domain: 82
```

Note: The default configuration for the G8264CS is the following:

```
RS G8264CS# show boot fabric
Current fabric settings:
Fabric           : ON
Switch Mode      : Member
Fabric Trunk Ports: 1:17-1:35
Fabric VLAN      : 4090
SI Fabric domain: 1
```

Only one of the G8264CS switches should be designated as **Master** in the **Switch Mode** field of the **Saved fabric settings** section. If neither or both of the G8264CS switches are designated as the master switch:

- a. Use the following commands on one G8264CS switch to designate it as the master node candidate.

```
RS G8264CS# boot fabric mode master
```

- b. Use the following command on the other G8264CS switch to designate it as the member node candidate.

```
RS G8264CS# boot fabric mode member
```

5. Reboot each G8264CS switch.


```
RS G8264CS# reload
WARNING: The running-config is different to startup-config.
Confirm operation without saving running-config to startup-config
(y/n) ? y
```

Note: All configuration changes using boot commands are automatically saved and are not show in user configuration dumps.

Connecting and Attaching the SI Fabric Devices

Once both G8264CS switches have completely rebooted, the SI Fabric components (G8264CS and SI4093 switches) can be connected together.

Figure 3. Simplified SI Fabric



Consider the example configuration in [Figure 3](#). Two G8264CS switches are used (one as a master node, and one as a backup). Two paired SI4093 devices in one Flex System chassis are shown, though up to 18 SI4093 switches in up to 9 chassis are supported.

1. Connect the appropriate Ethernet cables between participating devices:
 - Connect the two G8264CS switches together.
 - Connect the master G8264CS switch to each participating SI4093 or EN4093R switch.
 - Connect the backup G8264CS switch to each participating SI4093 or EN4093R switch.

In this example, two ports are used between each device for robust availability. However, devices can be connected using from 1 to 8 ports each.

Note: The port numbers are shown as symmetrical: each G8264CS switch uses the same port numbers to connect to each other, and also to any specific SI4093 or EN4093R. However, this is not a requirement; participating devices can be connected using any available fabric ports.

2. At the master G8264CS, verify that the participating devices are detected:

```

RS G8264CS> enable
RS G8264CS# show fabric attached-switches
Attached Switches in Fabric:
asnum      UUID                      Bay      MAC                      csnum      State      Type
-----
A1         N/A                      N/A      74:99:75:20:0e:00      C1         IN_FABRIC G8264CS
A2         669d8dccfb95492bb727d18313c82289 3      08:17:f4:5f:40:00      ATTACH     EN4093R
A3         7fb9a95c6a484e0296c418cae2bbfe5a 2      08:17:f4:84:3b:00      ATTACH     SI4093
A4         c6b9ceff4440486cb65bd1e2e3fe3a94 1      08:17:f4:84:3d:00      ATTACH     EN4093R
A5         c6b9ceff4440486cb65bd1e2e3fe3a94 4      08:17:f4:84:40:00      ATTACH     SI4093
A6         c6b9ceff4440486cb65bd1e2e3fe3a94 3      08:17:f4:a7:55:00      ATTACH     SI4093
A7         N/A                      N/A      08:17:f4:fb:cc:00      C2         IN_FABRIC G8264CS
A8         c6b9ceff4440486cb65bd1e2e3fe3a94 2      34:40:b5:73:92:00      ATTACH     EN4093R
A9         7fb9a95c6a484e0296c418cae2bbfe5a 4      34:40:b5:73:9c:00      ATTACH     SI4093
A10        7fb9a95c6a484e0296c418cae2bbfe5a 3      74:99:75:00:d1:00      ATTACH     SI4093
A11        b88989be60824783af6956a441cd0c61 3      74:99:75:1c:6b:00      ATTACH     EN4093R
A12        7fb9a95c6a484e0296c418cae2bbfe5a 1      74:99:75:1c:7d:00      ATTACH     EN4093R
A13        669d8dccfb95492bb727d18313c82289 2      74:99:75:5d:e1:00      ATTACH     SI4093
A14        669d8dccfb95492bb727d18313c82289 1      74:99:75:7f:7d:00      ATTACH     SI4093

```

Devices listed in the IN_FABRIC state are fully participating in the SI Fabric. At this initial stage, only the G8264CS devices will be IN_FABRIC. Flex switches will be in the ATTACH state, indicating that they have been detected but have not yet been configured to join SI Fabric operation. As no traffic is forwarded, internal ports for Flex switches in the ATTACH state are automatically disabled.

Each detected switch is automatically assigned a unique, non-configurable Attached Switch Number (asnum), based on the order in which it joins the SI Fabric. In addition, a logical identifier known as the Configured Switch Number (csnum) should be configured. During initial setup, the csnum values for master and backup G8264CS switches will be automatically assigned as C1 and C2 respectively, and csnum values for each Flex switch will be unassigned.

Configuring and Binding the SI4093/EN4093R Switches

Newly detected Flex switches are initially isolated from SI Fabric operation. Until they are explicitly configured to operate within the SI Fabric, they will appear in the ATTACH state and their internal ports will remain automatically disabled.

To approve the Flex switches for SI Fabric operation, each member Flex switch listed in the ATTACH state must be configured with a `csnum`. Once the `csnum` is assigned, the Flex switch will transition to the `IN_FABRIC` state and it will receive its configuration from the Master.

All Flex switch configuration is performed at the master G8264CS switch. On the master G8264CS switch, the Flex `csnum` values may be configured automatically for all switches at once, or manually for each individual switch.

Note: Fabric bind command is non-disruptive. When a new switch is added to the fabric, the command binds it to the next available switch number (`csnum`).

- To automatically assign `csnum` values for all attached Flex switches, use the following commands:

```
RS G8264CS# configure terminal
RS G8264CS(config)# fabric bind
```

- To manually assign a desired `csnum` value for any specific Flex switch, use the following configuration mode command:

```
RS G8264CS(config)# fabric switch-number <csnum> bind <asnum>
```

Repeat the manual command for each participating Flex switch.

Note: The master G8264CS always has a `csnum` of `C1`, and the backup G8264CS always has a `csnum` of `C2`. These cannot be manually changed.

When a new `csnum` is assigned, the member switch will leave and rejoin the SI Fabric and update its configuration accordingly. The following sequence occurs:

- If the desired `csnum` currently belongs to another member switch, that member switch will be placed in the ATTACH state and wait for a new `csnum` to be assigned before it may rejoin the SI Fabric.
- The `csnum` and binding assignment will be saved in the SI Fabric configuration as individual `csnum` entries with the physical switch MAC address (for G8264CS) or `csnum` and Flex System UUID plus IO bay number (for SI4093 or EN4093R).

In future reboots, each SI Fabric switch will retain the `csnum` saved in its binding configuration. The `asnum` values, however, may change, depending on the specific sequence in which switches are returned to service.

You can use the show fabric attached-switches command to verify that the participating devices are properly configured. For example:

```

RS G8264CS(config)# show fabric attached-switches
Attached Switches in Fabric:
asnum      UUID                               Bay      MAC                               csnum     State      Type
-----
A1         N/A                               N/A      74:99:75:20:0e:00                C1        IN_FABRIC  G8264CS
A2         669d8dccb95492bb727d18313c82289 3        08:17:f4:5f:40:00                C9        IN_FABRIC  EN4093R
A3         7fb9a95c6a484e0296c418cae2bbfe5a 2        08:17:f4:84:3b:00                C12       IN_FABRIC  SI4093
A4         c6b9ceff4440486cb65bd1e2e3fe3a94 1        08:17:f4:84:3d:00                C3        IN_FABRIC  EN4093R
A5         c6b9ceff4440486cb65bd1e2e3fe3a94 4        08:17:f4:84:40:00                C6        IN_FABRIC  SI4093
A6         c6b9ceff4440486cb65bd1e2e3fe3a94 3        08:17:f4:a7:55:00                C5        IN_FABRIC  SI4093
A7         N/A                               N/A      08:17:f4:fb:cc:00                C2        IN_FABRIC  G8264CS
A8         c6b9ceff4440486cb65bd1e2e3fe3a94 2        34:40:b5:73:92:00                C4        IN_FABRIC  EN4093R
A9         7fb9a95c6a484e0296c418cae2bbfe5a 4        34:40:b5:73:9c:00                C14       IN_FABRIC  SI4093
A10        7fb9a95c6a484e0296c418cae2bbfe5a 3        74:99:75:00:d1:00                C13       IN_FABRIC  SI4093
A11        b88989be60824783af6956a441cd0c61 3        74:99:75:1c:6b:00                C15       IN_FABRIC  EN4093R
A12        7fb9a95c6a484e0296c418cae2bbfe5a 1        74:99:75:1c:7d:00                C11       IN_FABRIC  EN4093R
A13        669d8dccb95492bb727d18313c82289 2        74:99:75:5d:e1:00                C8        IN_FABRIC  SI4093
A14        669d8dccb95492bb727d18313c82289 1        74:99:75:7f:7d:00                C7        IN_FABRIC  SI4093

```

The example output shows all detected switches configured with csnum values operating in the IN_FABRIC state. This indicates that the binding process is complete and all devices in the SI Fabric are acting as a single logical switch.

Reverting SI Fabric Members to Standalone Operation

You can remove any member switch from the SI Fabric and revert it to act as a standalone switch.

At the master G8264CS switch, use the following commands to revert any member switch and re-install a standalone software image:

```
RS G8264CS> enable
RS G8264CS# copy tftp {image1|image2|boot-image} address <IP address>
filename <standalone image filename> {data-port|mgt-port} switch-number <asnum>
```

Configuring the SI Fabric Management IP Addresses and Gateway

To provide out-of-band IP Management and SNMP connection to the SI Fabric, configure the an IP address and IP gateway at the master G8264CS switch.

1. View the attached switches MAC addresses and other information.

```
RS G8264CS> enable
RS G8264CS# show fabric attached-switches
Attached Switches in Fabric:
asnum      UUID                               Bay      MAC                               csnum     State      Type
-----
A1         N/A                               N/A      74:99:75:20:0e:00                C1        IN_FABRIC G8264CS
A2         669d8dccb95492bb727d18313c82289 3        08:17:f4:5f:40:00                C9        IN_FABRIC EN4093R
A3         7fbaa95c6a484e0296c418cae2bbfe5a 2        08:17:f4:84:3b:00                C12       IN_FABRIC SI4093
A4         c6b9ceff4440486cb65bd1e2e3fe3a94 1        08:17:f4:84:3d:00                C3        IN_FABRIC EN4093R
A5         c6b9ceff4440486cb65bd1e2e3fe3a94 4        08:17:f4:84:40:00                C6        IN_FABRIC SI4093
A6         c6b9ceff4440486cb65bd1e2e3fe3a94 3        08:17:f4:a7:55:00                C5        IN_FABRIC SI4093
A7         N/A                               N/A      08:17:f4:fb:cc:00                C2        IN_FABRIC G8264CS
A8         c6b9ceff4440486cb65bd1e2e3fe3a94 2        34:40:b5:73:92:00                C4        IN_FABRIC EN4093R
A9         7fbaa95c6a484e0296c418cae2bbfe5a 4        34:40:b5:73:9c:00                C14       IN_FABRIC SI4093
A10        7fbaa95c6a484e0296c418cae2bbfe5a 3        74:99:75:00:d1:00                C13       IN_FABRIC SI4093
A11        b88989be60824783af6956a441cd0c61 3        74:99:75:1c:6b:00                C15       IN_FABRIC EN4093R
A12        7fbaa95c6a484e0296c418cae2bbfe5a 1        74:99:75:1c:7d:00                C11       IN_FABRIC EN4093R
A13        669d8dccb95492bb727d18313c82289 2        74:99:75:5d:e1:00                C8        IN_FABRIC SI4093
A14        669d8dccb95492bb727d18313c82289 1        74:99:75:7f:7d:00                C7        IN_FABRIC SI4093
```

2. Assign the master G8264CS switch a management IP address, based on its MAC address.

```
RS G8264CS# configure terminal
RS G8264CS(config)# interface ip mgmt ip
RS G8264CS(config-if-mgmt)# mac a8:97:dc:05:70:00 ip address
172.16.50.24 255.255.0.0 enable
```

3. Assign the backup G8264CS switch a management IP address, based on its MAC address.

```
RS G8264CS(config-if-mgmt)# mac a8:97:dc:10:04:00 ip address
172.16.50.23 255.255.0.0 enable
```

4. Assign an IP gateway to the master and backup nodes, based on their MAC addresses:

```
RS G8264CS(config-if-mgmt)# exit
RS G8264CS(config)# ip gateway mgmt ip mac a8:97:dc:05:70:00 address
172.16.1.1 enable
RS G8264CS(config)# ip gateway mgmt ip mac a8:97:dc:10:04:00 address
172.16.1.1 enable
```

5. Verify the configuration.

```
RS G8264CS(config)# show running-config
Current configuration:
!
version "8.3"
switch-type "Lenovo Networking Operating System RackSwitch G8264CS,
SI
Fabric"
iscli-new
fabric switch-number 1 mac a8:97:dc:05:70:00
fabric switch-number 2 mac a8:97:dc:10:04:00
!
interface ip mgmt ip
mac a8:97:dc:05:70:00 ip address 172.16.50.24
mac a8:97:dc:05:70:00 ip netmask 255.255.0.0
mac a8:97:dc:05:70:00 enable
!
mac a8:97:dc:10:04:00 ip address 172.16.50.23
mac a8:97:dc:10:04:00 ip netmask 255.255.0.0
mac a8:97:dc:10:04:00 enable
!
ip gateway mgmt ip mac a8:97:dc:05:70:00 address 172.16.1.1
ip gateway mgmt ip mac a8:97:dc:05:70:00 enable
!
ip gateway mgmt ip mac a8:97:dc:10:04:00 address 172.16.1.1
ip gateway mgmt ip mac a8:97:dc:10:04:00 enable
!
end
```

In-band IP Management capability is also available for the SI Fabric logical switch with two Data IP interfaces over data ports. IP address and additional Telnet and SNMP access can be configured via the current Master node of the logical switch.

6. Assign and verify the first In-band management IP address and IP gateway via master node for SI Fabric management.

```
RS G8264CS(config)# interface ip data 1
RS G8264CS(config-ip-if-data)# ip address 172.16.60.24 255.255.0.0
enable
Warning: BOOTP will be disabled
WARNING: IP interface DATA1 and IP interface MGT are on the same
subnet.
WARNING: MGT IP interface (MAC: a8:97:dc:05:70:00) and DATA1 IP
interface have overlapped subnets.
WARNING: MGT IP interface (MAC: a8:97:dc:10:04:00) and DATA1 IP
interface have overlapped subnets.
...
RS G8264CS(config-ip-if-data)# show interface ip
Interface information:
DATA1: IP4 172.16.60.24      255.255.0.0      172.16.255.255,  vlan 1,
UP
MGT:   IP4 172.16.50.23     255.255.255.128  10.241.33.127,  vlan
4095, UP
No DHCP info available
[floating]: IP4 10.241.33.8      255.255.255.128  10.241.33.127
```

Floating IP Address

Lenovo SI Fabric also supports floating IP which can be used for out-of-band management of fabric by connecting to the current Master (regardless of whether the current Master is the configured Master or the Backup).

1. Disable DHCP on Master and Backup:

```
RS G8264CS(config)# no system dhcp mac <master_mac_addr>
RS G8264CS(config)# no system dhcp mac <backup_mac_addr>
```

2. Configure Management IP address on Master and Backup interfaces:

```
RS G8264CS(config)# interface ip mgmt ip
RS G8264CS(config-ip-if-mgmt)# mac <master_mac_addr> ip address <ip_addr>
<netmask>
RS G8264CS(config-ip-if-mgmt)# mac <backup_mac_addr> ip address <ip_addr>
<netmask>
RS G8264CS(config-ip-if-mgmt)#exit
```

3. Configure gateways on Master and Backup:

```
RS G8264CS(config)# ip gateway mgmt ip mac <master_mac_addr> address
<gateway> enable
RS G8264CS(config)# ip gateway mgmt ip mac <backup_mac_addr> address
<gateway> enable
```

4. Configure the floating IP:

```
RS G8264CS(config)# interface ip mgmt ip
RS G8264CS(config-ip-if-mgmt)# floating ip address <ip_addr> <netmask>
RS G8264CS(config-ip-if-mgmt)#exit
```

Note: Master and Backup IP interfaces settings are minimal for configuring floating IP.

SNMP, Telnet, and SSH Access

Many SI Fabric features may be managed remotely using SNMP (and Switch Center), Telnet, or SSH.

Remote management requires prior configuration of the SI Fabric Management IP address and gateway (see [“Configuring the SI Fabric Management IP Addresses and Gateway”](#) on page 28).

SNMP and Switch Center Access

The SI Fabric can be managed via Simple Network Management Protocol (SNMP) using an SNMP management tool such as Switch Center. SNMP resource management and switch access are similar to operating a standalone G8264CS switch.

SI Fabric SNMP settings must be configured prior to use. To configure and verify SNMP settings, use the following configuration commands at the Master G8264CS switch:

```
RS G8264CS(config)# snmp-server version v1v2v3
RS G8264CS(config)# snmp-server read-community "public"
RS G8264CS(config)# snmp-server write-community "private"

RS G8264CS(config)# show snmp-server
Current SNMP params:
  Read community string: "public"
  Write community string: "private"
  SNMP state machine timeout: 5 minutes
  Trap source address: mgmt
  Authentication traps disabled.
  All link up/down traps enabled.

Current SNMP trap hosts:

Current v1/v2 access enabled
```

You can view SNMP-server counters as follows.

```
RS G8264CS(config)# show snmp-server counters
-----
SNMP statistics:
snmpInPkts:                317028   snmpInBadVersions:          0
snmpInBadC'tyNames:        2         snmpInBadC'tyUses:          0
snmpInASNParseErrs:        0         snmpEnableAuthTraps:       2
snmpOutPkts:                317022   snmpInBadTypes:             0
snmpInTooBig:               0         snmpInNoSuchNames:         0
snmpInBadValues:           0         snmpInReadOnlys:           0
snmpInGenErrs:              0         snmpInTotalReqVars:        1413078
snmpInTotalSetVars:         0         snmpInGetRequests:         463
snmpInGetNexts:             306916   snmpInSetRequests:          0
snmpInGetResponses:        0         snmpInTraps:                0
snmpOutTooBig:              0         snmpOutNoSuchNames:        0
snmpOutBadValues:          0         snmpOutReadOnlys:          0
snmpOutGenErrs:             0         snmpOutGetRequests:        0
snmpOutGetNexts:            0         snmpOutSetRequests:        0
snmpOutGetResponses:       307379   snmpOutTraps:               0
snmpSilentDrops:            0         snmpProxyDrops:            0
```

Telnet Access

Telnet allows remote access to the SI Fabric CLI. By default, Telnet access to the SI Fabric is disabled.

SI Fabric access and resource management via Telnet are similar to operating a standalone G8264CS switch.

To enable and verify Telnet settings, user the following commands at the Master G8264 switch:

```
RS G8264CS(config)# access telnet enable
RS G8264CS(config)# show access
Current System Access settings:

IP Management currently allowed from *ALL* IP addresses

Usernames:
  user   - enabled   - offline
  oper   - disabled  - offline
  admin  - Always Enabled - online   1 session.
Current User ID table:

Current strong password settings:
  strong password status: disabled

SNMP access currently read-write
Telnet access currently enabled on TCP port 23
TFTP occurs over port 69
```

SSH Access

Secure Shell (SSH) allows for more secure remote access to the SI Fabric CLI than simple Telnet provides. SSH is enabled by default.

SI Fabric access and resource management via SSH are similar to operating a standalone G8264CS switch.

Displaying SI Fabric Configuration

Display the current running configuration to verify SI Fabric settings:

```
RS G8264CS(config)# show running-config
Current configuration:
!
version "8.2.1.0"
switch-type "Lenovo Networking Operating System RackSwitch G8264CS, SI
Fabric"
iscli-new
fabric switch-number 1 mac a8:97:dc:05:70:00
fabric switch-number 2 mac a8:97:dc:10:04:00
fabric switch-number 3 universal-unic-id 63ceda2ec8cc46bf9eb8b3c14b2cefco
bay 1
fabric switch-number 4 universal-unic-id 63ceda2ec8cc46bf9eb8b3c14b2cefco
bay 2
!
!
snmp-server version v1v2v3
!
snmp-server read-community "public"
snmp-server write-community "private"
!
no system bootp
!
access telnet enable
!
interface ip mgmt ip
    mac a8:97:dc:05:70:00 ip address 172.16.50.24
    mac a8:97:dc:05:70:00 ip netmask 255.255.0.0
    mac a8:97:dc:05:70:00 enable
!
    mac a8:97:dc:10:04:00 ip address 172.16.50.23
    mac a8:97:dc:10:04:00 ip netmask 255.255.0.0
    mac a8:97:dc:10:04:00 enable
!
ip gateway mgmt ip mac a8:97:dc:05:70:00 address 172.16.1.1
ip gateway mgmt ip mac a8:97:dc:05:70:00 enable
!
ip gateway mgmt ip mac a8:97:dc:10:04:00 address 172.16.1.1
ip gateway mgmt ip mac a8:97:dc:10:04:00 enable
!
end
```

Chapter 3. Managing the Flex SI Fabric Solution

After you have performed the Flex System Interconnect Fabric (SI Fabric) initial configuration steps (see [Chapter 2, "Initial Configuration"](#)), complete the relevant management actions:

- ["SI4093 and EN4093R License Activation Keys"](#) on page 36
- ["Managing Systems through Master and Backup G8264CS Switches"](#) on page 38

SI4093 and EN4093R License Activation Keys

License keys determine the number of available ports on the Flex switches:

- Basic License

The basic license is active by default on the Flex switch. It provides the use of 24 ports as follows:

- Internal ports INTA1 through INTA14
- External ports EXT1 through EXT10

- Optional Upgrade License 1

In addition to all ports provided by the basic license, the following are available:

- Internal ports INTB1 through INTB14
- Two External QSFP+ ports. By default, each is configured for operation as four 10 Gbps Ethernet uplinks (EXT15 through EXT22) using the appropriate system divider cables. However, if high-bandwidth connections are preferred, each QSFP+ port can independently configured to operate in undivided 40 Gbps Ethernet mode.

- Optional Upgrade License 2

Upgrade license 2 requires that upgrade license 1 be installed first. In addition to all ports provided by the basic license and upgrade license 1, the following are available:

- Internal ports INTC1 through INTC14
- External SFP+ ports EXT11 through Port EXT14

Additionally, the user has the option to select any ports as long as the total bandwidth for all the enabled ports does not exceed the current installed license. For further details, see [“Flexible Port Mapping” on page 39](#).

Obtaining Activation Keys

The upgrade licenses can be acquired using the *Lenovo System x Features on Demand* (FoD) website:

<http://www.ibm.com/systems/x/fod/>

You can also use the website to review and manage licenses, and to obtain additional help, if required.

Note: A Lenovo ID and password are required to log into the FoD website. If you do not yet have a Lenovo ID, you can register at the website.

Activation keys are provided as files that must be uploaded to the Flex switch. To acquire an activation key, use the FoD website to purchase an Authorization Code. You will need to provide the unique ID (UID) of the specific Flex switch where the key will be installed. The UID is the last 12 characters of the Flex switch serial number. This serial number is located on the Part Number (PN) label and is also displayed during successful login to the device.

Download the activation key file from the FoD site when available in your account.

Installing Activation Keys

Once FoD activation key files have been acquired, they must be installed on the Flex switch. The example below depicts use of the Master Command Line Interface (CLI), but other device interfaces (such as SNMP) may also be used.

When installing licenses, please note the following requirements:

- The Flex switch must be reset to activate any newly installed licenses.
- The 64 Port License (Upgrade 2) will not function unless the 46 Port License (Upgrade 1) is also present. If installing both upgrades at the same time, upload both keys prior to resetting the Flex switch.

To install activation keys, complete the following steps:

1. Log in to the master G8264CS.
2. At the CLI prompt, enter the following commands:

```
RS G8264CS> enable
RS G8264CS# configure terminal
RS G8264CS(config)# software-key
RS G8264CS(Software-Key)# enakey address <server IP address> keyfile <key
filename> protocol <tftp or sftp> switch <3-20> mgt-port | data-port
```

3. Follow the prompts to enter the appropriate parameters, including the file transfer protocol and server parameters.

Note: Repeat the enakey command for any additional keys being installed.

Transferring Activation Keys

Licenses keys are based on the unique Flex switch device serial number and are non-transferable.

In the event that the Flex switch must be replaced, a new activation key must be acquired and installed. When the replacement is handled through Lenovo Service and Support, your original license will be transferred to the serial number of the replacement unit and you will be provided a new license key.

Managing Systems through Master and Backup G8264CS Switches

The SI Fabric is designed to enable master node and backup node between the G8264CS switches pair to provide continuous availability. All SI Fabric management is done using the master G8264CS switch. If the current master node fails, the current backup G8264CS switch will detect the absence of the master node and become the new master node of the SI Fabric. See [“Configuring the SI Fabric Management IP Addresses and Gateway”](#) on page 28 for details.

The Management IP address of the new master node is used to continue the resource management of the SI Fabric. If the data IP interface is used, the new master node will automatically assume the responsibility of resource management.

Note: If both the master and backup G8264CS switches fail, all Member Flex switch will also reset with all Internal server ports in disabled state and wait to rejoin the SI Fabric logical switch once the master G8264CS node recovers.

Deploying Staggered Image Upgrades for Master G8264CS Switch

The SI Fabric has a staggered mode for image upgrades. This staggered image upgrade provides network redundancy for the compute node, with redundant connection to its local Flex switches. All servers are still reachable during the image upgrade.

As with the traditional image upgrade process, the staggered upgrade downloads the SI Fabric boot image and OS software image to master node, and lets the master node push the new images to all member switches.

The staggered approach differs from the traditional image upgrade process in that instead of requiring a simultaneous reboot of the entire logical switch, the staggered update permits the master to reboot switches on a sequential basis.

Following is the staggered image upgrade sequence:

- After the new boot image and OS image are downloaded and installed on all switches, the master node will reboot the backup node. The master node waits for the backup node to rejoin before it reboots itself to preserve failover capabilities to the current backup node.
- After both new master and new backup nodes are booted with new image, the new master node will reboot all Flex switches in I/O module bay 1 of all Flex System chassis.
- After all rebooted Flex switches rejoin, the new master node will reboot all Flex switches in IO Module bay 2 of all Flex System chassis. This rolling reboot will take place at all Flex switches in I/O module bay 3 then at all Flex switches in I/O Module bay 4.

Note: Boot image should be downloaded and installed, prior to downloading OS image via staggered upgrade.

During the staggered upgrade process, the compute node with redundant connection will ensure that at least one server link is available and prevents any server connectivity interruption.

Deploying Staggered Image Reboot for Master G8264CS Switch

The SI Fabric allows the user to reboot the switch (without performing an image upgrade) and still have server access.

Staggered reload can be executed only if the new boot image is identical with the current one. All switches in the SI Fabric need to run the same software version.

Flexible Port Mapping

The SI Fabric allows administrators to manually enable or disable specific switch ports within the limitations of the installed licenses' bandwidth. Each switch in the SI Fabric that supports licensing has bandwidth restrictions determined by its license level.

Commands associated with flexible port mapping can only be run from the master switch and can have an additional parameter.

To add or remove ports of a SI Fabric switch to/from the port map by specifying the switch's configured number and port number, use the following command:

```
RS G8264CS(config)# [no] boot port-map <ports alias or member>
```

To reset the port map configuration to the default settings, use the following command:

```
RS G8264CS(config)# default boot port-map [<csnum>]
```

If a configured switch number is specified, the command will reset the port map configuration only for the selected switch.

Microburst Detection

Microbursts are short peaks in data traffic that manifest as a sudden increase in the number of data packets transmitted over a specific millisecond-level time frame, potentially overwhelming network buffers. Microburst detection allows users to analyze and mitigate microburst-related incidents, thus preventing network congestion.

To enable or disable microburst detection, use the following command:

```
RS G8264CS(config)# [no] microburst enable
```

By default, microburst detection is disabled.

Connecting Uplink to the Data Center Network

The SI Fabric provides uplink connection from both G8264CS switches to the data center network. The uplink port group feature provides layer-2 loop-free connectivity without the complexity of STP operation. The uplink connection can be a single physical interface, a static portchannel trunk group, a static LACP trunk group, or a hotlink pair.

When the SI Fabric is initially configured, all Ethernet ports that are not configured as SI Fabric ports at either G8264CS switches can be used as uplink connection. Initially, these Ethernet ports for uplink connection are placed in the Black-hole VLAN, which prevents any network loop. While an uplink port is in Black-hole VLAN, the port will not receive/forward traffic.

When an uplink connection is assigned to a target data VLAN or to a set of target data VLANs from black-hole VLAN (VLAN 4091 in this example), an uplink group is formed.

The uplink group prevents any ingress traffic from looping back to data center network and allows the SI Fabric to operate without use of Spanning Tree Protocol (STP).

When an uplink connection is removed from its last data VLAN, the SI Fabric moves this uplink connection back to the Black-hole VLAN and ensures no network loop.

The example below shows the commands and commented text for the uplink connection creation and assignments at the SI Fabric switches in IO Module bay 1 and bay 2 of the same Flex System chassis.

1. An uplink group is formed when the administrator configures an uplink connection to a target data VLAN or a set of target data VLANs from the Black-hole VLAN. The uplink group in the SI Fabric prevents other uplink connection from adding to this target data VLAN(s).
2. Configure interfaces and assign their data VLANs.

You can use a mix of one or more single physical interfaces, static portchannel trunk groups, LACP trunk groups, and hotlink pairs as uplink connections. For example:

- A single physical interface:

```
RS G8264CS(config)# interface port 1:2
RS G8264CS(config-if)# switchport access vlan 100
RS G8264CS(config-if)# exit
```

- Static portchannel:

```
RS G8264CS(config)# portchannel 288 port 1:3,2:3 enable
RS G8264CS(config)# interface port 1:3,2:3
RS G8264CS(config-if)# switchport access vlan 101
RS G8264CS(config-if)# exit
```


o LACP:

```
RS G8264CS(config)# portchannel 289 lACP key 2890
RS G8264CS(config)# interface port 1:4,2:4
RS G8264CS(config-if)# lACP key 2890
RS G8264CS(config-if)# lACP mode active
RS G8264CS(config-if)# switchport access vlan 102
RS G8264CS(config-if)# exit
```

3. Configure hotlinks for uplink redundancy:

a. Configure hotlink triggers:

```
RS G8264CS(config)# hotlinks trigger 1 enable
RS G8264CS(config)# hotlinks trigger 1 master port 1:5
RS G8264CS(config)# hotlinks trigger 1 backup port 2:5
```

b. Globally enable the hotlinks feature:

```
RS G8264CS(config)# hotlinks ena
```

c. Add the trigger interfaces to their uplink VLANs.

If the trigger interface is composed of access ports:

```
RS G8264CS(config)# interface port 1:5,2:5
RS G8264CS(config-if)# switchport access vlan 104
RS G8264CS(config-if)# exit
```

Otherwise, if the trigger interface is composed of ports in trunk mode:

```
RS G8264CS(config)# interface port 1:5,2:5
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan 104
RS G8264CS(config-if)# exit
```

4. Verify port configuration:

```
RS G8264CS(config)# show interface info 1:2-1:5,2:1-2:5
```

Defining Server Ports

In full-fabric mode, multiple uplink groups can be connected to servers via uplink ports. Once defined as server port, the port can become member of any VLANs.

Uplink groups server port loop detector is based on BPDU Guard Feature. If a port configured as server port receives a BPDU, the port is automatically error disabled.

To define a port as a server port use the following command:

```
RS G8264CS(config)# system server-ports port <ports alias or member>
```

Connecting Downlink Compute Nodes

The SI Fabric provides downlink connections from the Flex switches' internal ports-only compute nodes. Factory default settings configure all Internal ports on Flex switches for downlink connection in a default VLAN.

Note: If you assign downlink connections, the SI Fabric does not impose the uplink connection requirement mentioned previously.

The SI Fabric also supports the Unified Fabric Port (UFP) protocol between internal ports as downlink connections and server NICs. The downlink connection provisioning at the SI Fabric is similar to the common network provisioning at a standalone Flex switch.

In order to maintain system redundancy, two Flex switches are used, installed in neighboring bays in the Flex System chassis: either Bay 1 and Bay 2, or alternately in Bay 3 and Bay 4. When you configure the Flex downlink connection to a compute node that has redundant NICs or NIC teaming, be sure that the switches and NICs are in the same chassis, and that the same internal port number is used in both paired Flex switches.

The example below shows the downlink connection creation and assignments at the SI Fabric logical switch made up by two G8264CS switches and two Flex switches in IO Module bay 1 and bay 2 of the same Flex System chassis.

1. Confirm that the paired Flex switches are in the same Flex System chassis and in neighboring IO module bays.

```
RS G8264CS(config)# show fabric attached-switches
Attached Switches in Fabric:
asnum      UUID                               Bay      MAC                               csnum     State      Type
-----
A1         N/A                               N/A      74:99:75:20:0e:00                C1        IN_FABRIC G8264CS
A2         669d8dccfb95492bb727d18313c82289 3        08:17:f4:5f:40:00                C9        IN_FABRIC EN4093R
A3         7fbaa95c6a484e0296c418cae2bbfe5a 2        08:17:f4:84:3b:00                C12       IN_FABRIC SI4093
A4         c6b9ceff4440486cb65bd1e2e3fe3a94 1        08:17:f4:84:3d:00                C3        IN_FABRIC EN4093R
A5         c6b9ceff4440486cb65bd1e2e3fe3a94 4        08:17:f4:84:40:00                C6        IN_FABRIC SI4093
A6         c6b9ceff4440486cb65bd1e2e3fe3a94 3        08:17:f4:a7:55:00                C5        IN_FABRIC SI4093
A7         N/A                               N/A      08:17:f4:fb:cc:00                C2        IN_FABRIC G8264CS
A8         c6b9ceff4440486cb65bd1e2e3fe3a94 2        34:40:b5:73:92:00                C4        IN_FABRIC EN4093R
A9         7fbaa95c6a484e0296c418cae2bbfe5a 4        34:40:b5:73:9c:00                C14       IN_FABRIC SI4093
A10        7fbaa95c6a484e0296c418cae2bbfe5a 3        74:99:75:00:d1:00                C13       IN_FABRIC SI4093
A11        b88989be60824783af6956a441cd0c61 3        74:99:75:1c:6b:00                C15       IN_FABRIC EN4093R
A12        7fbaa95c6a484e0296c418cae2bbfe5a 1        74:99:75:1c:7d:00                C11       IN_FABRIC EN4093R
A13        669d8dccfb95492bb727d18313c82289 2        74:99:75:5d:e1:00                C8        IN_FABRIC SI4093
A14        669d8dccfb95492bb727d18313c82289 1        74:99:75:7f:7d:00                C7        IN_FABRIC SI4093
```

The example information shows that the Flex switches with csnum C3 and C4 are in the same Flex System (their UUID is the same) in neighboring IO Module bays: Bay 1 and Bay 2 respectively.

2. View default VLAN information.

```
RS G8264CS(config)# show vlan 1
VLAN          Name                Status          Ports
-----
1             Default VLAN        ena             1:17-1:35 2:17-2:35
                                     3:INTA1-3:EXT22
                                     4:INTA1-4:EXT22
...
```

In this example, the default VLAN 1 information includes all possible downlink ports from all Flex member switches.

3. Configure the downlink connection in pairs.

```
RS G8264CS(config)# interface port 3:INTA1,4:INTA1
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan 1060
```

In this example, the pair of Internal ports 3:INTA1 and 4:INTA1 represent physical internal ports on two members Flex switch and are connected to the compute node in Bay 1 of the Flex System chassis.

4. Verify the interface configuration.

```
RS G8264CS(config-if)# show interface info 3:INTA1,4:INTA1
Alias  Port Tag   Type   RMON Lrn PVID  DESCRIPTION
VLAN(s)
      Trk           NVLAN
-----
3:INTA1 257  y  Internal  d  e   1  3:INTA1      1 1060
4:INTA1 385  y  Internal  d  e   1  4:INTA1      1 1060
RS G8264CS(config-if)# exit
```

This example shows the available downlink ports on the C3 and C4 Flex switches and their assigned VLANs.

Connecting to Storage Resources

The SI Fabric facilitates integrated end-to-end FC/FCoE storage connectivity. In this configuration, G8264CS switches act as FCoE gateways for bridging FCoE and Fibre Channel networks, and embedded Flex switches as FCoE aggregators. It connects FCoE enabled compute nodes and storage nodes with Converged Enhanced Ethernet (CEE) support.

SI Fabric serves storage connectivity requirements in three different ways - FIP Snooping Bridge, NPV Gateway mode, and Full fabric mode. The SI Fabric in FIP Snooping configuration connects to an external FCoE Forwarder (FCF) to provide secure FCoE device connectivity. In NPV Gateway mode, the SI Fabric can provide connectivity to existing SAN fabric via Omni ports (1:53-1:64 and 2:53-2:64) on G8264CS as uplinks. In Full fabric mode, the SI Fabric serves as FC-BB-5 compliant FC/FCoE SAN which provides connectivity to FC/FCoE server/storage.

FIP Snooping Bridge

Follow these steps to enable FCoE function at the SI Fabric, comprised of two G8264CS switches and two Flex switches.

1. Enable FCoE function by turning on CEE and global FIP Snooping.

```
RS G8264CS(config)# cee enable  
RS G8264CS(config)# fcoe fips enable
```

2. Create a typical FCoE VLAN 1002 for internal FCoE nodes and external FCF. if required.

```
RS G8264CS(config)# vlan 1002  
VLAN 1002 is created.
```

3. Place the internal FCoE port and external FCF port into the FCoE VLAN with VLAN trunking enabled.


```
RS G8264CS(config)# interface port 1:20,3:INTA5  
RS G8264CS(config-if)# switchport mode trunk  
RS G8264CS(config-if)# switchport trunk allowed vlan 1,1002  
RS G8264CS(config-if)# exit
```

4. Keep default FCoE settings (FIP Snooping/FCF mode) on FCoE/FCF ports.

```
RS G8264CS(config)# show fcoe fips port 1:20,3:INTA5  
  
Port 1:20 FIP Snooping Configuration:  
FIP snooping: enabled, FCF mode: auto  
  
Port 3:INTA5 FIP Snooping Configuration:  
FIP snooping: enabled, FCF mode: auto  
...
```

NPV Gateway Mode

Figure 4. NPV Gateway example



1. Enable NPV Gateway function starting with enabling each pair of Omni ports with Fibre Channel mode at both G8264CS switches (which automatically enables VLAN trunking).

```
RS G8264CS(config)# system port 1:53-1:64,2:53-2:64 type fc
```

2. Create an FC VLAN for NPV Gateway functionality on first G8264CS and add at least one FC port from it as an uplink to SAN fabric. It is recommended to have more than one uplink to SAN fabric. Each G8264CS switch can accommodate up to 12 NPV VLANs. To have traffic redundancy, configure different NPV VLAN from the other G8264CS. Also specify the FC port(s), which has the external connection to upstream full-fabric SAN switch under each NPV enabled VLAN.

```
RS G8264CS(config)# vlan 1060
VLAN 1060 is created.

RS G8264CS(config-vlan)# interface fc 1:57,1:58
RS G8264CS(config-fc)# switchport trunk allowed vlan 1,1060

RS G8264CS(config-fc)# interface port 3:INTA1,3:INTA2
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan add 1060
RS G8264CS(config-if)# exit

RS G8264CS(config)# vlan 1060
RS G8264CS(config-vlan)# npv enable
RS G8264CS(config-vlan)# npv traffic-map external-interface 1:57,1:58
RS G8264CS(config-vlan)# exit
```

3. Configure another NPV enabled VLAN for second G8264CS switch.

```
RS G8264CS(config)# vlan 2060
VLAN 1060 is created.

RS G8264CS(config-vlan)# interface fc 2:57,2:58
RS G8264CS(config-fc)# switchport trunk allowed vlan 1,2060

RS G8264CS(config-fc)# interface port 4LINTA1,4:INTA2
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan add 2060
RS G8264CS(config-if)# exit

RS G8264CS(config)# vlan 2060
RS G8264CS(config-vlan)# npv enable
RS G8264CS(config-vlan)# npv traffic-map external-interface 2:57,2:58
RS G8264CS(config-vlan)# exit
```

4. Verify NPV status.

```
RS G8264CS(config)# show npv status
Unit:1 VLAN: 1060      NPV enabled
Unit:2 VLAN: 2060      NPV enabled

RS G8264CS(config)# show npv traffic-map
-----
      VLAN      Source Ports      NP-Uplink Dest Ports
-----
      1060      empty             1:57,1:58
      2060      empty             2:57,2:58
```

Full-Fabric Mode

In full-fabric mode, the SI Fabric supports E_Port feature for up to four FC ISL links between G8264CS switches to expand fabric connectivity.

Follow these steps to enable full-fabric function on the same SI Fabric:

1. Enable full-fabric FC/FCoE switch function by using these FC ports at either G8264CS switches for direct connectivity to FC/FCoE storage nodes.

```
RS G8264CS(config)# cee enable
RS G8264CS(config)# fcoe fips enable
RS G8264CS(config)# system port 1:53-1:64,2:53-2:64 type fc
```

2. Create one FC VLAN for full FC/FCoE network for the first G8264CS switch and add at least one FC port from it and the other FCoE internal ports. Each G8264CS switch can have only 1 full fabric VLAN. The full fabric VLAN should be a different VLAN from the other G8264CS switch for storage traffic redundancy.

```
RS G8264CS(config)# vlan 1024
VLAN 1024 is created.

RS G8264CS(config-vlan)# interface fc 1:55,1:56
RS G8264CS(config-fc)# switchport trunk allowed vlan 1,1024

RS G8264CS(config-fc)# interface port 3:INTA3,3:INTA4
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan add 1024
RS G8264CS(config-if)# exit

RS G8264CS(config)# vlan 1024
RS G8264CS(config-vlan)# fcf enable
RS G8264CS(config-vlan)# exit
```

3. Configure another full FC VLAN for the second G8264CS switch.

```
RS G8264CS(config)# vlan 2024
VLAN 1024 is created.

RS G8264CS(config-vlan)# interface fc 2:55,2:56
RS G8264CS(config-fc)# switchport trunk allowed vlan 1,2024

RS G8264CS(config-fc)# interface port 4:INTA3,4:INTA4
RS G8264CS(config-if)# switchport mode trunk
RS G8264CS(config-if)# switchport trunk allowed vlan add 2024
RS G8264CS(config-if)# exit

RS G8264CS(config)# vlan 2024
RS G8264CS(config-vlan)# fcf enable
```

4. Configure E_Ports.

E_ports (expansion ports) connect two full-fabric switches to form an inter-switch link (ISL). This is an optional configuration which is useful only when user wanted to expand the SAN fabric to more than one switch. Maximum of four full-fabric switches can be interconnected. Only Fibre Channel port types can be configured as E-ports. These ports must be members of a Fibre Channel VLAN. Use the following commands to configure E_ports:

```
RS G8264CS(config-fc)# interface fc 1:54  
RS G8264CS(config-fc)# switchport trunk allowed vlan add 2024  
RS G8264CS(config-fc)# type e  
  
RS G8264CS(config-fc)# exit
```

5. Configure FC zoning.

Each FC port can only be assigned to one FC VLAN and each FC VLAN can only include FC ports from the same G8264CS switch

```
RS G8264CS(config)# zone name master vlan 1024  
RS G8264CS(config-zone)# member pwwn 10:00:5c:f3:fc:6e:4a:a1  
RS G8264CS(config-zone)# member pwwn 50:05:07:68:05:04:09:2c  
RS G8264CS(config-zone)# zoneset name SC vlan 1024  
RS G8264CS(config-zoneset)# member master  
RS G8264CS(config-zoneset)# zoneset activate name SC vlan 1024  
RS G8264CS(config-zoneset)# exit  
  
RS G8264CS(config)# zone name backup vlan 2024  
RS G8264CS(config-zone)# member pwwn 50:05:07:68:05:08:09:2c  
RS G8264CS(config-zone)# member pwwn 10:00:5c:f3:fc:6e:4a:a5  
RS G8264CS(config-zone)# zoneset name SC2 vlan 2024  
RS G8264CS(config-zoneset)# member backup  
RS G8264CS(config-zoneset)# zoneset activate name SC2 vlan 2024  
RS G8264CS(config-zoneset)# exit
```

Chapter 4. Flex SI Fabric -Specific ISCLI Commands

The Flex System Interconnect Fabric (SI Fabric) requires a new group of ISCLI commands to manage the new set of capabilities. These commands come in addition to the existing SI4093, EN4093R and G8264CS commands.

For the existing SI4093, EN4093R and G8264CS commands, see the respective product ISCLI guides.

Information Commands

Command Syntax and Usage

show fabric

Displays the current SI Fabric information.

Command mode: All

show interface fabric-trunk <trunk ID> interface-counters

Displays the current SI Fabric trunk statistics.

Command mode: All

show system black-hole

Displays the current configured SI Fabric Black-hole VLAN.

Command mode: All

show system internal-VLAN

Displays the current SI Fabric internal-VLAN space.

Command mode: All

show system chassis swn <switch number>

Displays the current chassis related information at the specified switch.

Command mode: All

Configuration Commands

Command Syntax and Usage

system black-hole vlan <VLAN number>

Configures the SI Fabric black-hole VLAN

Command mode: Global configuration

system internal-VLAN <VLAN number or range> [**black-hole vlan** <VLAN number>]

Configures the SI Fabric internal VLAN or VLANs.

Command mode: Global configuration

copy tftp {image1|image2} **address** <IP address> **filename** <filename> {**data-port**|**mgt-port**} **staggered-upgrade** [**delay** <minutes (1–30)>]

Copies image from TFTP server and upgrade SI Fabric members in staggered mode.

Command mode: Global configuration

copy ftp {image1|image2} {**data-port**|**mgt-port**} **staggered-upgrade** [**delay** <minutes (1–30)>]

Copies image from FTP server and upgrade SI Fabric members in staggered mode.

Command mode: Global configuration

copy sftp {image1|image2} {**data-port**|**mgt-port**} **staggered-upgrade** [**delay** <minutes (1–30)>]

Copies image from SFTP server and upgrade SI Fabric members in staggered mode.

Command mode: Global configuration

reload staggered [**delay** <delay>]

Starts a rolling reload based on the time delay in minutes between switch reboots. The default value is one minute between switch reboots.

Command mode: Global configuration

fabric bind

Configure all attached switch units.

Command mode: Global configuration

[no] logging log {**all**|<feature>}

Displays a list of features for which syslog messages can be generated. You can choose to enable/disable specific features (such as vlans, stg, or ssh), or enable/disable syslog on all available features.

Command mode: Global configuration

copy log [**swn** <switch number>] {**stfp**|**tftp** [**address** <address>]} [**filename** <filename>]}

Copies syslog content to an external host using SFTP or TFTP.

Command mode: Global configuration

IP Interface Configuration Commands

You can use the Command Line Interface (CLI) for making, viewing, and saving IP interface configuration changes.

Command Syntax and Usage

interface ip data {1|2}

Enter Data IP interface mode.

Command mode: Global configuration

interface ip mgmt {ip|ipv6}

Enter Management IP interface mode.

Command mode: Global configuration

mac <MAC address> ip address <IP address> <IP netmask> enable

Configure the IP address using dotted decimal notation.

Command mode: Interface IP (Management)

mac <MAC address> ipv6 address <IPv6 address> <IPv6 prefix> enable

Configure the IPv6 address using hexadecimal format with colons.

Command mode: Interface IP (Management)

show interface ip [data|mgmt]

Displays the current IP interface settings.

Command mode: All

ip gateway data {ip|ipv6} address {<IPv4 address>|<IPv6 address>} enable

Configures the IP address of the default IP gateway over Data IP interface.

Command mode: Global configuration

ip gateway mgmt {ip|ipv6} mac <MAC address> address {<IPv4 address>|<IPv6 address>} enable

Configures the IP address of the default IP gateway over Management IP interface.

Command mode: Global configuration

show ip gateway {data|mgmt}

Displays the current IP interface settings.

Command mode: All

Software Key Commands

You can use the Command Line Interface (CLI) for updating and reviewing switch FoD software license.

Command Syntax and Usage

software-key

Enter software key command menu.

Command mode: Global configuration

enakey address <IP address> **keyfile** <filename> [**protocol** {tftp|sftp}] **switch** <switch number> {**data-port**|**mgt-port**}]

Enable software key.

Command mode: Software Key Menu

invkeys address <IP address> [**infile** <filename> **protocol** {tftp|sftp}] **switch** <switch number> {**data-port**|**mgt-port**}]

Upload inventory installed activation keys.

Command mode: Software Key Menu

ptkeys address <IP address> **key** {Upgrade1|Upgrade2} [**protocol** {tftp|sftp}] **file** <filename> **switch** <switch number> {**data-port**|**mgt-port**}]

Upload Software Key.

Command mode: Software Key Menu

rmkey key {Upgrade1|Upgrade2} **switch** <switch number>

Remove Software Key.

Command mode: Global configuration

show software-key [**switch** <switch number>]

Show software licensing keys.

Command mode: All

Boot Options

To use following new Boot Options commands, you must be logged in to the switch as the administrator. The new Boot Options commands provide options for the SI Fabric.

Command Syntax and Usage

boot fabric si-fabric-domain *<domain number>*

Configures the SI Fabric domain number on switches. The same domain number needs to be applied on both switches to form the SI Fabric.

Command mode: Global configuration

boot fabric fabric-trunk *<list of ports>*

Configures the ports used to connect the switch to the SI Fabric. Enter only 10Gb external ports.

Command mode: Global configuration

si-fabric sif-errup

Turns on the ports that were disabled (`errdis`) during SI Fabric merger.

Command mode: Global configuration

show boot staggered-upgrade

Show status of staggered upgrade of the member switches in the SI Fabric.

Command mode: All

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.

Use this information to 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 about obtaining service. IBM is Lenovo's preferred service provider for the System x, Flex System, and NeXtScale System products.

Before you call, make sure that you have taken these steps to try to solve the problem yourself.

If you believe that you require warranty service for your Lenovo product, the service technicians will be able to assist you more efficiently if you prepare before you call.

- 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 and firmware for the product (unless it is covered by an additional maintenance contract). Your service technician will request that you upgrade your software and firmware if the problem has a documented solution within a software 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.
- 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.
 - Hardware and Software Maintenance agreement contract numbers, if applicable
 - Machine type number (Lenovo 4-digit machine identifier)
 - Model number
 - Serial number
 - Current system UEFI and firmware levels
 - Other pertinent information such as error messages and logs

- Start the process of determining a solution to your problem by making the pertinent information available to the service technicians. The IBM service technicians can start working on your solution as soon as you have completed and submitted an Electronic Service Request.

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.