

IBM Flex System Interconnect Fabric



# Solution Guide

for IBM Networking OS 7.8



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

**First Edition, June 2014. Part Number 00CG967**

**© Copyright IBM Corporation 2014**

US Government Users Restricted Rights—Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.



---

# Contents

<b>Preface</b> . . . . .	<b>.7</b>
Who Should Use This Guide . . . . .	.7
Typographic Conventions . . . . .	.7
Additional Resources . . . . .	.8
<b>Chapter 1. Introduction</b> . . . . .	<b>.9</b>
SI Fabric Architecture . . . . .	.9
Point of Delivery . . . . .	10
Hyper Distributed Fabric Protocol . . . . .	10
SI Fabric Components . . . . .	11
SI4093 Switches . . . . .	11
G8264CS Switches . . . . .	11
SI Fabric Features . . . . .	11
Layer-2 Switching Capabilities . . . . .	11
Network Simplification . . . . .	11
Management Simplification . . . . .	12
Storage Integration . . . . .	12
Scalable Point of Delivery . . . . .	12
Limitations . . . . .	12
<b>Chapter 2. Initial Configuration</b> . . . . .	<b>13</b>
Hardware Planning . . . . .	13
Initial System Configuration Overview . . . . .	13
Before Configuring the SI Fabric . . . . .	14
Using the ISCLI . . . . .	14
Updating and Rebooting All Participating Switches . . . . .	14
Specifying SI Fabric Settings on the G8264CS . . . . .	16
Connecting and Attaching the SI Fabric Devices . . . . .	18
Configuring and Binding the SI4093 Switches . . . . .	19
Reverting SI Fabric Members to Stand-Alone Operation . . . . .	20
Configuring the SI Fabric Management IP Addresses and Gateway . . . . .	21
SNMP, Telnet, and SSH Access . . . . .	23
SNMP and Switch Center Access . . . . .	23
Telnet Access . . . . .	24
SSH Access . . . . .	24
Displaying SI Fabric Configuration . . . . .	25
<b>Chapter 3. Managing the Flex SI Fabric Solution</b> . . . . .	<b>27</b>
SI4093 License Activation Keys . . . . .	27
Obtaining Activation Keys . . . . .	27
Installing Activation Keys . . . . .	28
Transferring Activation Keys . . . . .	28
Managing Systems through Master and Backup G8264CS Switches . . . . .	29
Deploying Staggered Image Upgrades for Master G8264CS Switch . . . . .	29
Connecting Uplink to the Data Center Network . . . . .	30
Connecting Downlink Compute Nodes . . . . .	32
Connecting to Storage Resources . . . . .	33
E-Port Mode . . . . .	33
NPV Gateway Mode . . . . .	34
Enabling Full-Fabric Function . . . . .	36

<b>Chapter 4. Flex SI Fabric Solution-Specific ISCLI Commands . . . . .</b>	<b>39</b>
Information Commands . . . . .	39
Configuration Commands . . . . .	40
IP Interface Configuration Commands . . . . .	41
Software Key Commands . . . . .	42
Boot Options . . . . .	43
<b>Appendix A. Getting Help and Technical Assistance . . . . .</b>	<b>45</b>
Before You Call . . . . .	45
Using the Documentation . . . . .	45
Getting Help and Information on the World Wide Web . . . . .	46
Software Service and Support . . . . .	46
Hardware Service and Support . . . . .	46
IBM Taiwan product service . . . . .	47

---

## Preface

The *IBM Flex SI Fabric Solution Guide* describes how to configure and use the IBM Networking OS 7.8 software on the IBM System Networking RackSwitch G8264CS (referred to as G8264CS throughout this document) and the IBM Flex System Fabric SI4093 System Interconnect Module (referred to as SI4093) to achieve the IBM Flex System Interconnect Fabric (referred to as the SI Fabric).

For documentation about the physical installation of the SI4093 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>
<b>ABC123</b>	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	<code>host# sys</code>
<i>&lt;ABC123&gt;</i>	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 &lt;IP address&gt;</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>

Table 1. *Typographic Conventions (continued)*

Typeface or Symbol	Meaning	Example
{   }	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}
<b>AaBbCc123</b>	This block type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the <b>Save</b> button.

---

## Additional Resources

Additional information about installing and configuring the devices of the IBM Flex System Interconnect Fabric is available on the IBM support website:

<http://www.ibm.com/support>

<http://www.redbooks.ibm.com>

Components of the Flex SI Fabric Solution, and their documentation, includes:

- IBM Flex System Fabric SI4093 System Interconnect Module (SI4093)
  - SI4093 *User Guide* (for physical features and installation)
  - SI4093 *Application Guide for IBM Networking OS 7.8*
  - SI4093 *ISCLI Reference for IBM Networking OS 7.8*
- IBM System Networking RackSwitch G8264CS
  - G8264CS *Installation Guide* (for physical features and installation)
  - G8264CS *Application Guide for IBM Networking OS 7.8*
  - G8264CS *ISCLI Reference for IBM Networking OS 7.8*

The product support website also provides access to help, service, and technical assistance.

If further help is needed, the warranty card received with your products provide details for contacting a customer support representative. If you are unable to locate this information, please contact your reseller. Before you call, prepare the following information:

- Serial number of the product
- Software release version number
- Brief description of the problem and the steps you have already taken
- Technical support dump information (**show tech-support**)



# Chapter 1. Introduction

This *Solution Guide* introduces the IBM 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 Fiber 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:

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

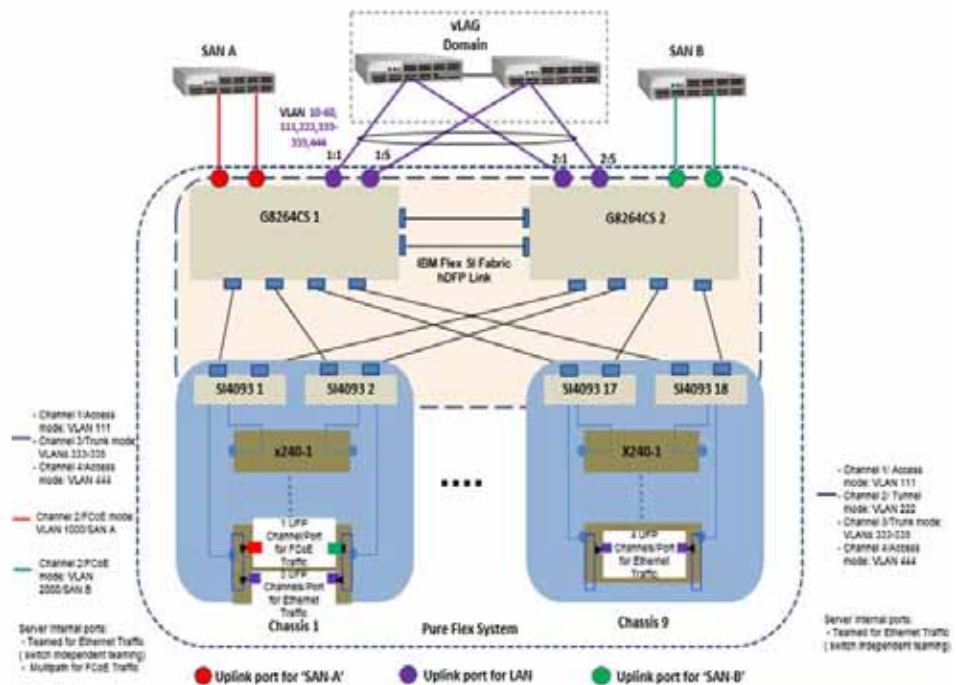
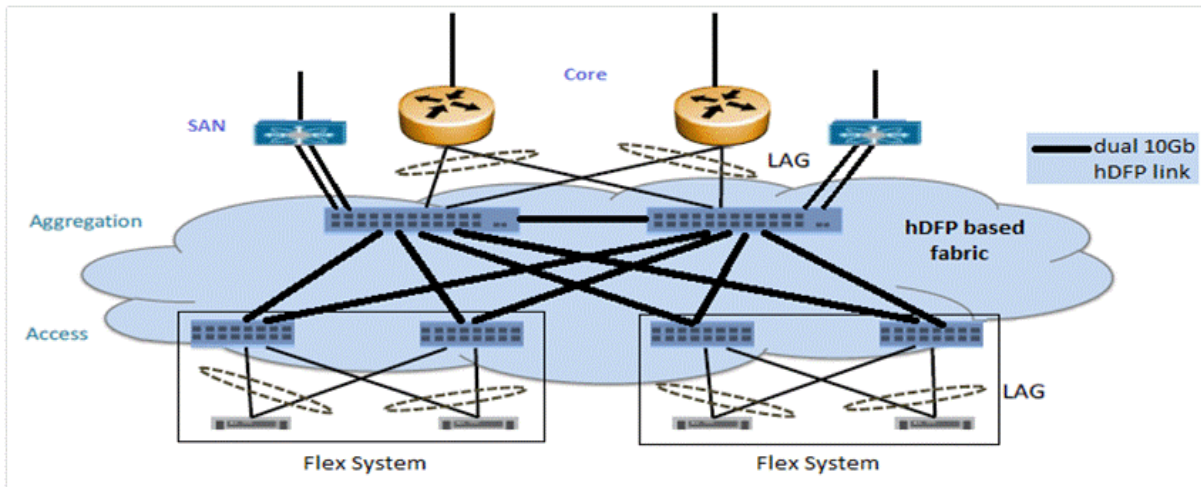


Figure 2 on page 10 illustrates the SI Fabric connections to storage and computing resources.

Figure 2. Flex SI Fabric Solution 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 IBM 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 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 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, G8264CS, and storage hardware components.

### SI4093 Switches

The IBM Flex System Fabric SI4093 System Interconnect Module (SI4093) provides a transparent network interface between server elements within the chassis and the upstream data and storage networking infrastructure. The standard SI4093 provides 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 provides 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 IBM 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

### 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

- Integrates Fiber Channel Forwarder (FCF) to provide end-to-end Fiber Channel over Ethernet (FCoE) storage functionality within the POD without the need of expensive Fiber Channel switch
- Supports single-fabric mode topology, dual-fabric mode topology, and full-fabric topology

## Management Simplification

The SI Fabric provides the following management simplification features:

- High availability with master and backup nodes
- Upgrade with no service down time
- 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 Flex System V7000 storage node and access to external SAN storage infrastructure, such as IBM Storwize V7000.

## 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 SI Fabric Solution limitations for the current release:

- The SI Fabric requires IBM Networking OS version 7.7.5 or greater for all stand-alone switches.
- IBM does not support mixing 10Gb and 40Gb SI Fabric port speeds in SI Fabric logical switches. IBM recommends that you reserve additional SI Fabric ports on both G8264CS switches for future expansion.
- All external ports on SI4093 devices are automatically configured as Fabric Ports, requiring no additional configuration. Omni Ports on the RackSwitch G8264CS cannot be configured as Fabric Ports.

---

## Chapter 2. Initial Configuration

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

---

### Hardware Planning

The minimum configuration requires two G8264CS switches, two Flex System chassis populated with two SI4093 switches at each chassis, two half-width Flex System compute nodes with 10Gb CNAs at each chassis, and Storwize V7000 storage system.

The SI Fabric supports up to two G8264CS switches, up to nine Flex System chassis populated with two SI4093 switches at each chassis, up to 126 half-width Flex System compute nodes and external Storwize V7000 storage system.

The SI Fabric acts as a simple pass-thru switch that provides a transparent network interface between server elements within the chassis and the upstream data and storage networking infrastructure.

- The base SI4093 switch provides 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.
- The SI4093 switch in the SI Fabric provides the convergence of traditional Ethernet traffic with Fibre Channel over Ethernet (FCoE) storage.

---

### Initial System Configuration Overview

To deploy the Flex SI Fabric Solution from a stand-alone topology:

1. Upgrade and reboot all stand-alone G8264CS switches and SI4093 switches with SI Fabric software images (see [“Updating and Rebooting All Participating Switches” on page 14](#)).
2. 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 16](#)).
3. Connect two 10Gb Ethernet cables as SI Fabric links between the two G8264CS switches. Also connect two 10Gb Ethernet cables from each SI4093 switch to each G8264CS switch. See [“Connecting and Attaching the SI Fabric Devices” on page 18](#).
4. Once the master G8264CS detects all attached SI4093 switches, bind them to the SI Fabric (see [“Configuring and Binding the SI4093 Switches” on page 19](#)).
5. 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 21](#)).
6. Enable additional management access (see [“SNMP, Telnet, and SSH Access” on page 23](#)).
7. If required, install FoD license keys on SI4093 switches (see [“SI4093 License Activation Keys” on page 27](#)).

## 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

## 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 or Telnet, 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.

---

## Updating and Rebooting All Participating Switches

Before stand-alone G8264CS and SI4093 switches can participate in the SI Fabric, each must have the SI Fabric boot and software images installed.

**Note:** SI Fabric images can be installed only on switches running IBM Networking OS (N/OS) version 7.7.5 or later. If your switches are running any earlier version of N/OS, update them to 7.7.5 or later prior to installing SI Fabric software.

Use the following procedure to update N/OS 7.7.5 (or later) stand-alone G8264CS and SI4093 switches to use SI Fabric 7.8 software:

1. Ensure that all stand-alone switches are running software version IBM Networking OS 7.7.5 or greater. Use the `show version` and `show boot` commands:

```
RS G8264CS> enable
RS G8264CS# show version | include Version
Software Version 7.7.5 (FLASH image2), active configuration.

RS G8264CS# show boot
Currently set to boot software image2, active config block.
...
```

2. Place the SI Fabric software images on a TFTP server or an FTP server for download, and configure all stand-alone switches to use this TFTP server or FTP server as the image server.
3. Enable the IP address and management port.

```
RS G8264CS# configure terminal
RS G8264CS(config)# interface ip mgmt ip
RS G8264CS(config-ip-if-mgmt)# mac 74:99:75:7f:77:00 ip address 172.16.2.214
255.255.255.0 enable
RS G8264CS(config-ip-if-mgmt)# exit
RS G8264CS(config)# ip gateway mgmt ip mac 74:99:75:7f:77:00 address 172.16.1.1
ena
Feb 24 23:01:07 G8264CS-R NOTICE ip: MGT port default gateway 172.16.1.1
operational
```

4. Use ping to ensure you can connect the switch to the server where the software is stored.

```
RS G8264CS(config)# ping 172.16.101.11 2 mgt-port
Connecting via MGT port.
[host 172.16.101.11, max tries 2, delay 1000 msec, length 0, ping source N/S,
ttl 255, tos 0]
172.16.101.11: #1 ok, RTT 1 msec.
172.16.101.11: #2 ok, RTT 1 msec.
Ping finished.
```

5. Before starting the download:

- If the switch is running N/OS 7.7.5, use the following command to set the switch to accept SI Fabric software images:

```
RS G8264CS(config)# copy hfdp-mode
```

- If the switch with N/OS 7.8 or later software release, use the following command to set the switch to accept SI Fabric software images:

```
RS G8264CS(config)# copy si-fabric-mode
```

6. Download and install both the Boot image and OS image for the SI Fabric on all stand-alone G8264CS and SI4093 switches which are intended for use in the SI Fabric. For example:

```
RS G8264CS(config)# copy tftp boot-image address 172.16.101.11 filename
G8264-CS-GbFSIM-10G-FSIF-7.8.2.0_Boot.img mgt-port

boot kernel currently contains Software Version 7.7.5
New download will replace boot kernel with file
"G8264-CS-GbFSIM-10G-FSIF-7.8.2.0_Boot.img"
  from SFTP/FTP/TFTP server 172.16.101.11.
Connecting via MGT port.
Confirm download operation (y/n) ? y
Image download started to boot kernel
File appears valid
Download in progress
...
```

**Note:** You must download the SI Fabric software image to the same image location (image1 or image2) on all G8264CS and SI4093 switches intended for use in the SI Fabric.

7. Reboot all G8264CS and SI4093 switches.

After the switches are rebooted, all switches boot with factory-default configuration and blank primary and backup configurations.

8. After you have downloaded and rebooted all SI Fabric components, use `show boot` to confirm that the new SI Fabric software image is used.

```
RS G8264CS> enable
RS G8264CS# show boot
Current running image version: 7.8.2.0
Currently set to boot software image2, active config block.
NetBoot: disabled, NetBoot tftp server: , NetBoot cfgfile:
Current FLASH software:
  image1: version 7.7.5.0, downloaded 13:29:09 Mon Mar 17, 2014
           NormalPanel, Mode SI Fabric
  image2: version 7.8.2.0, downloaded 19:48:13 Thu Mar 27, 2014
           NormalPanel, Mode SI Fabric
  boot kernel: version 7.8.2.0, Mode SI Fabric
Currently scheduled reboot time: none

Copy Mode: SI Fabric
```

In this example, `Mode SI Fabric` indicates where SI Fabric software is used.

Once participating GS8264CS and SI4093 switches have been rebooted with SI Fabric software, no further configuration is required through the SI4093 switch CLI. All SI Fabric configuration (including any SI4093 settings), will be performed through the master G8264CS switch connection established using the following procedures.

---

## Specifying SI Fabric Settings on the G8264CS

After you have rebooted all G8264CS and SI4093 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:** IBM 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 43](#) 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
```

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 switch.
  - Connect the backup G8264CS switch to each participating SI4093 switch.

In this example, two ports are used between each device for robust availability. However, devices can be connected using from 1 to 6 ports each.

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

- 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    a8:97:dc:05:70:00 C1     IN_FABRIC G8264CS
A2  63ceda2ec8cc46bf9eb8b3c14b2cefc0 1    34:40:b5:4f:53:00   ATTACH SI4093
A3  63ceda2ec8cc46bf9eb8b3c14b2cefc0 2    34:40:b5:50:90:00   ATTACH SI4093
A4             N/A              N/A    a8:97:dc:10:04:00   C2     IN_FABRIC G8264CS

```

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`. SI4093 switches will be in the `ATTACH` state, indicating that they have been detected but have not yet been approved to join SI Fabric operation. For security, internal ports for SI4093 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 configurable logical identifier known as the Configured Switch Number (`csnum`) is available. 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 SI4093 switch will be unassigned.

---

## Configuring and Binding the SI4093 Switches

For security, newly detected SI4093 switches are initially isolated from SI Fabric operation. Until they are explicitly permitted to operate within the SI Fabric, they will appear in the `ATTACH` state and their internal ports will remain automatically disabled.

To approve the SI4093 switches for SI Fabric operation, each member SI4093 switch listed in the `ATTACH` state must be configured with a `csnum`. Once the `csnum` is assigned, the SI4093 switch will transition to the `IN_FABRIC` state.

All SI4093 configuration is performed at the master G8264CS switch. On the master G8264CS switch, the SI4093 `csnum` values may be configured automatically for all switches at once, or manually for each individual switch.

- To automatically assign `csnum` values for all attached SI4093 switches, use the following commands:

```

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

```

- To manually assign a desired `csnum` value for any specific SI4093 switch, use the following configuration mode command:

```

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

```

Repeat the manual command for each participating SI4093 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 port list 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 or `csnum` and Flex System UUID plus IO bay number.

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      a8:97:dc:05:70:00 C1     IN_FABRIC G8264CS
A2  63ceda2ec8cc46bf9eb8b3c14b2cefc0  1       34:40:b5:4f:53:00 C3     IN_FABRIC SI4093
A3  63ceda2ec8cc46bf9eb8b3c14b2cefc0  2       34:40:b5:50:90:00 C4     IN_FABRIC SI4093
A4             N/A                               N/A      a8:97:dc:10:04:00 C2     IN_FABRIC G8264CS
```

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 Stand-Alone Operation

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

At the master G8263CS switch, use the following commands to revert any member switch and re-install a stand-alone software image:

```
RS G8264CS> enable
RS G8264CS# copy non-si-fabric-mode
RS G8264CS# copy tftp {image1|image2} address <IP address> filename
<stand-alone 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    a8:97:dc:05:70:00 C1 IN_FABRIC G8264CS
A2  63ceda2ec8cc46bf9eb8b3c14b2cefc0 1  34:40:b5:4f:53:00 C3 IN_FABRIC SI4093
A3  63ceda2ec8cc46bf9eb8b3c14b2cefc0 2  34:40:b5:50:90:00 C4 IN_FABRIC SI4093
A4             N/A              N/A    a8:97:dc:10:04:00 C2 IN_FABRIC G8264CS
```

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-ip-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-ip-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-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 "7.8.0.50"
switch-type "IBM 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, DOWN
MGT:   IP4 172.16.50.23    255.255.0.0    172.16.255.255,  vlan 4095, UP
No DHCP info available

MAC: a8:97:dc:05:70:00
MGT:   IP4 172.16.50.24    255.255.0.0    172.16.255.255,  vlan 4095, up
No DHCP info available
...
```

---

## 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 21).

## 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 stand-alone 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
snmpInTooBigs:      0       snmpInNoSuchNames:      0
snmpInBadValues:    0       snmpInReadOnlys:        0
snmpInGenErrs:      0       snmpInTotalReqVars:     1413078
snmpInTotalSetVars: 0       snmpInGetRequests:      463
snmpInGetNexts:     306916  snmpInSetRequests:      0
snmpInGetResponses: 0       snmpInTraps:             0
snmpOutTooBigs:     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 stand-alone G8264CS switch.

To enable and verify Telnet settings, use 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 stand-alone 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 "7.8.0.50"
switch-type "IBM 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 63ceda2ec8cc46bf9eb8b3c14b2cefc0 bay 1
fabric switch-number 4 universal-unic-id 63ceda2ec8cc46bf9eb8b3c14b2cefc0 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 SI Fabric Solution (SI Fabric) initial configuration steps (see [Chapter 2, “Initial Configuration”](#)), complete the relevant management actions:

- [“SI4093 License Activation Keys” on page 27](#)
- [“Managing Systems through Master and Backup G8264CS Switches” on page 29](#)

---

### SI4093 License Activation Keys

License keys determine the number of available ports on the SI4093 switches:

- Basic License

The basic license is active by default on the SI4093. 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

### Obtaining Activation Keys

The upgrade licenses can be acquired using the *IBM 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:** An IBM ID and password are required to log into the FoD website. If you do not yet have an IBM ID, you can register at the website.

Activation keys are provided as files that must be uploaded to the SI4093. 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 SI4093 where the key will be installed. The UID is the last 12 characters of the SI4093 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 SI4093. The example below depicts use of the SI4093 Command Line Interface (CLI), but other device interfaces (such as SNMP) may also be used.

When installing licenses, please note the following requirements:

- The SI4093 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 SI4093.

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 <keyfilename>
```

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.

4. Once the key file has been uploaded, reset the target SI4093 switch or switches to activate any newly installed license:

```
RS G8264CS(Software-Key)# reload switch <csnum>
```

The system prompts you to confirm your request. Once confirmed, the specified `csnum` will reboot with the new licenses.

## Transferring Activation Keys

Licenses keys are based on the unique SI0493 device serial number and are non-transferable.

In the event that the SI4093 must be replaced, a new activation key must be acquired and installed. When the replacement is handled through IBM 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 21](#) 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 SI4093 switches 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 SI4093 switches.

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 SI4093 switches in I/O module bay 1 of all Flex System chassis.
- After all rebooted SI4093 switches rejoin, the new master node will reboot all SI4093 switches in IO Module bay 2 of all Flex System chassis. This rolling reboot will take place at all SI4093 switches in I/O module bay 3 then at all SI4093 switches in I/O Module bay 4.

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.

---

## Connecting Uplink to the Data Center Network

The SI Fabric provides uplink connection from both G8264CS switches to the data center network. The dual uplink 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.

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

The SI Fabric, using the uplink group, stops other uplink connection from adding to this target data VLANs. 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
```

– LACP:

```
RS G8264CS(config)# portchannel 289 lACP key 2890 suspend-individual
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 trunks:

```
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
```

---

## Connecting Downlink Compute Nodes

The SI Fabric provides downlink connections from the SI4093 switches' internal ports-only compute nodes. Factory default settings configure all Internal ports on SI4093 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 stand-alone SI4093 switch.

In order to maintain system redundancy, two SI4093 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 SI4093 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 SI4093 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 SI4093 switches in IO Module bay 1 and bay 2 of the same Flex System chassis.

1. Confirm that the paired SI4093 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    74:99:75:7f:7a:00  C1     IN_FABRIC  G8264CS
A2  db8cf9c5719a4853b715f9cba8e3bab0  1     74:99:75:05:e7:00  C3     IN_FABRIC  SI4093
A3  db8cf9c5719a4853b715f9cba8e3bab0  2     74:99:75:d6:81:00  C4     IN_FABRIC  SI4093
A4             N/A                      N/A    74:99:75:7f:f7:00  C2     IN_FABRIC  G8264CS
...
```

The example information shows that the SI4093 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 SI4093 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 add 1060
```

In this example, the pair of Internal ports 3:INTA1 and 4:INTA1 represent physical internal ports on a member SI4093 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
-----
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 SI4093 switches. This indicates the C3 switch has both FOD license Upgrade 1 and Upgrade 2 applied with 42 Internal ports available, while C4 switch has default license with 14 Internal ports available.

---

## Connecting to Storage Resources

The SI Fabric provides integrated end-to-end FCoE storage connection with G864CS switches as FCoE gateways. This configuration bridges FCoE and Fibre Channel networks and SI4093 switches as FCoE aggregators. It connects FCoE enabled compute nodes and storage nodes with Converged Enhanced Ethernet (CEE) support.

With the Omni ports on G864CS switches, the SI Fabric can also provide native Fibre Channel (FC) connectivity to connect to existing SANs via port 1:53-1:64 and 2:53-2:64. The SI Fabric can serve as an NPV Gateway, which can uplink multiple Fibre Channel nodes to an existing full-fabric SAN switch. It can also serve as FC-BB-5 compliant full-fabric FC/FCoE SAN switch to connect to NPV Gateways or interconnecting Fibre Channel nodes.

## E-Port 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 FCoE function at the SI Fabric, comprised of two G8264CS switches and two SI4093 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 add 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 one FC VLAN for NPV Gateway function for first G8264CS switch and add at least one FC port from it and other FCoE internal ports - each G8264CS switch can have up to 11 NPV enabled VLAN and should be different VLANs from the other G8264CS switch for storage traffic redundancy. 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 add 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 add 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
```

## Enabling Full-Fabric Function

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

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 up to 1 full FC VLAN. The full FC 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 add 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 add 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.

In this example, add an FC port to the second FC VLAN and set it as an E\_Port. Also use one FC port from the first FC VLAN as E\_Port to build the FC ISL between G8264CS switches.

```
RS G8264CS(config-vlan)# interface fc 2:54
RS G8264CS(config-fc)# switchport trunk allowed vlan add 2024
RS G8264CS(config-fc)# type e

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 Solution-Specific ISCLI Commands

The Flex SI Fabric Solution (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 and G8264CS commands.

For the existing SI4093 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 si-fabric-mode**

Sets image copy mode to copy SI Fabric software image to member switches.

**Command mode:** Global configuration

---

**copy non-si-fabric-mode**

Sets image copy mode to copy non-SI Fabric stand-alone software image to member switches.

**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

---

**fabric bind**

Configure all attached switch units.

**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> [**invfile** <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 IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

---

### Before You Call

Before you call, make sure that you have taken these steps to try to solve 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.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM *Documentation CD* that comes with your system.
- Go to the IBM support website at <http://www.ibm.com/systems/support/> to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the documentation that is provided with your IBM product. The documentation that comes with IBM systems also describes the diagnostic tests that you can perform. Most systems, operating systems, and programs come with documentation that 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.

---

### Using the Documentation

Information about your IBM system and pre-installed software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, online documents, ReadMe files, and Help files.

See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software.

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/systems/support/> and follow the instructions. Also, some documents are available through the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

---

## Getting Help and Information on the World Wide Web

The IBM website has up-to-date information about IBM systems, optional devices, services, and support.

IBM System x<sup>®</sup> and xSeries<sup>®</sup> information is available at:

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

IBM BladeCenter information is available at:

<http://www.ibm.com/systems/bladecenter/>

IBM IntelliStation<sup>®</sup> information is available at:

<http://www.ibm.com/intellistation/>

You can find service information for IBM systems and optional devices at:

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

---

## Software Service and Support

Through the IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with System x and x Series servers, BladeCenter products, IntelliStation workstations, and appliances.

For information about which products are supported by Support Line in your country or region, see:

<http://www.ibm.com/services/sl/products/>

For more information about Support Line and other IBM services, see:

<http://www.ibm.com/services/>

For support telephone numbers, see:

<http://www.ibm.com/planetwide/>

In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

---

## Hardware Service and Support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to <http://www.ibm.com/partnerworld/> and click **Find Business Partners** on the right side of the page. For IBM support telephone numbers, see <http://www.ibm.com/planetwide/>. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

---

## IBM Taiwan product service

台灣 IBM 產品服務聯絡方式：  
台灣國際商業機器股份有限公司  
台北市松仁路7號3樓  
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation  
3F, No 7, Song Ren Rd.  
Taipei, Taiwan  
Telephone: 0800-016-888



Part Number: 00CG967

Printed in USA

(IP) P/N: 00CG967