

Flex System Fabric SI4093 System Interconnect Module

Product Guide (withdrawn product)

The Flex System™ Fabric SI4093 System Interconnect Module enables simplified integration of Flex System into your existing networking infrastructure and provides the capability of building simple connectivity for points of delivery (PODs) or clusters up to 252 nodes. The SI4093 requires no management for most data center environments, eliminating the need to configure each networking device or individual ports, thus reducing the number of management points. It provides a low latency, loop-free interface that does not rely upon spanning tree protocols, thus removing one of the greatest deployment and management complexities of a traditional switch. The SI4093 offers administrators a simplified deployment experience while maintaining the performance of intra-chassis connectivity. The SI4093 System Interconnect Module is shown in Figure 1.



Figure 1. Flex System Fabric SI4093 System Interconnect Module

Did you know?

The base switch configuration comes standard with 24x 10 GbE port licenses that can be assigned to internal connections or external SFP+ or QSFP+ ports with flexible port mapping. For example, this feature allows you to trade off four 10 GbE ports for one 40 GbE port (or vice versa) or trade off one external 10 GbE SFP+ port for one internal 10 GbE port (or vice versa). You then have the flexibility of turning on more ports when you need them using Features on Demand upgrade licensing capabilities that provide “pay as you grow” scalability without the need to buy additional hardware.

The SI4093 can be used in the Flex System Interconnect Fabric solution that reduces networking management complexity without compromising performance by lowering the number of devices that need to be managed by 95% (managing one device instead of 20). Interconnect Fabric simplifies POD integration into an upstream network by transparently interconnecting hosts to a data center network and representing the POD as a large compute element isolating the POD’s internal connectivity topology and protocols from the rest of the network.

Part number information

The SI4093 module is initially licensed for 24x 10 GbE ports. Further ports can be enabled with Upgrade 1 and Upgrade 2 license options. Upgrade 1 must be applied before Upgrade 2 can be applied. Table 1 shows the part numbers for ordering the module and the upgrades.

Table 1. Part numbers and feature codes for ordering

Description	Part number	Feature code
Interconnect module		
Flex System Fabric SI4093 System Interconnect Module	95Y3313	A45T
Features on Demand upgrades		
Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)	95Y3318	A45U
Flex System Fabric SI4093 System Interconnect Module (Upgrade 2)	95Y3320	A45V

The base part number for the interconnect module includes the following items:

- One Flex System Fabric SI4093 System Interconnect Module
- Important Notices Flyer
- Warranty Flyer
- Documentation CD-ROM

Note: SFP and SFP+ (small form-factor pluggable plus) transceivers or cables are not included with the module. They must be ordered separately (see Table 2).

The interconnect module does not include a serial management cable. However, Flex System Management Serial Access Cable, 90Y9338, is supported and contains two cables, a mini-USB-to-RJ45 serial cable and a mini-USB-to-DB9 serial cable, either of which can be used to connect to the interconnect module locally for configuration tasks and firmware updates.

The part numbers for the upgrades, 95Y3318 and 95Y3320, include the following items:

- Features on Demand Activation Flyer
- Upgrade authorization letter

The base module and upgrades are as follows:

- 95Y3313 is the part number for the base module, and it comes with 14 internal 10 GbE ports enabled (one to each compute node) and ten external 10 GbE ports enabled.
- 95Y3318 (Upgrade 1) can be applied on the base module when you take full advantage of four-port adapter cards installed in each compute node. This upgrade enables 14 additional internal ports, for a total of 28 ports. The upgrade also enables two 40 GbE external ports. This upgrade requires the base module.
- 95Y3320 (Upgrade 2) can be applied on top of the Upgrade 1 when you need more external bandwidth on the module or if you need additional internal bandwidth to the compute nodes with the six-port capable adapter cards. The upgrade will enable the remaining four external 10 GbE ports, plus 14 additional internal 10 GbE ports, for a total of 42 internal ports (three to each compute node).

Flexible port mapping: With Networking OS version 7.8 or later clients have more flexibility in assigning ports that they have licensed on the SI4093 which can help eliminate or postpone the need to purchase upgrades. While the base model and upgrades still activate specific ports, flexible port mapping provides clients with the capability of reassigning ports as needed by moving internal and external 10 GbE ports, or trading off four 10 GbE ports for the use of an external 40 GbE port. This is very valuable when you consider the flexibility with the base license and with Upgrade 1.

With flexible port mapping, clients have licenses for a specific number of ports:

- 95Y3313 is the part number for the base module, and it provides 24x 10 GbE ports licenses that can enable any combination of internal and external 10 GbE ports and external 40 GbE ports (with the use of four 10 GbE port licenses per one 40 GbE port).
- 95Y3318 (Upgrade 1) upgrades the base module by activation of 14 internal 10 GbE ports and two external 40 GbE ports which is equivalent to adding 22 more 10 GbE port licenses for a total of 46x 10 GbE port licenses. Any combination of internal and external 10 GbE ports and external 40 GbE ports (with the use of four 10 GbE port licenses per one 40 GbE port) can be enabled with this upgrade. This upgrade requires the base module.
- 95Y3320 (Upgrade 2) requires the base module and Upgrade 1 already be activated and simply activates all the ports on the SI4093 which is 42 internal 10 GbE ports, 14 external SFP+ ports, and two external QSFP+ ports.

Note: When both Upgrade 1 and Upgrade 2 are activated, flexible port mapping is no longer used because all the ports on the SI4093 are enabled.

Table 2 lists supported port combinations on the SI4093 and required upgrades.

Table 2. Supported port combinations (Part 1: Default port mapping)

Supported port combinations (Default port mapping)	Quantity required		
	Base module, 95Y3313	Upgrade 1, 95Y3318	Upgrade 2, 95Y3320
<ul style="list-style-type: none"> • 14x internal 10 GbE ports • 10x external 10 GbE ports 	1	0	0
<ul style="list-style-type: none"> • 28x internal 10 GbE ports • 10x external 10 GbE ports • 2x external 40 GbE ports 	1	1	0
<ul style="list-style-type: none"> • 42x internal 10 GbE ports† • 14x external 10 GbE ports • 2x external 40 GbE ports 	1	1	1

† This configuration leverages six of the eight ports on the CN4058 adapter available for IBM Power Systems™ compute nodes.

Table 2. Supported port combinations (Part 2: Flexible port mapping*)

Supported port combinations (Flexible port mapping)	Quantity required		
	Base module, 95Y3313	Upgrade 1, 95Y3318	Upgrade 2, 95Y3320**
<ul style="list-style-type: none"> • 24x 10 GbE ports (internal and external) or • 20x 10 GbE ports (internal and external) • 1x external 40 GbE ports or • 16x 10 GbE ports (internal and external) • 2x external 40 GbE ports 	1	0	0
<ul style="list-style-type: none"> • 46x 10 GbE ports (internal and external) or • 42x 10 GbE ports (internal and external) • 1x external 40 GbE ports or • 38x 10 GbE ports (internal and external) • 2x external 40 GbE ports 	1	1	0

* Flexible port mapping is available in Networking OS 7.8 or later.

** Upgrade 2 is not used with flexible port mapping because with Upgrade 2 all ports on the module become licensed and there is no need to reassign ports.

Supported cables and transceivers

With the flexibility of the SI4093 module, clients can take advantage of the technologies that they require for multiple environments:

- For 1 GbE links, clients can use RJ-45 SFP transceivers with UTP cables up to 100 m. Clients that need longer distances can use a 1000BASE-SX transceiver, which can drive distances up to 220 meters by using 62.5 μ multi-mode fiber and up to 550 meters with 50 μ multi-mode fiber, or the 1000BASE-LX transceivers that support distances up to 10 kilometers using single-mode fiber (1310 nm).
- For 10 GbE (on external SFP+ ports), clients can use SFP+ direct-attached copper (DAC) cables for in-rack cabling and distances up to 7 m. These DAC cables have SFP+ connectors on each end, and they do not need separate transceivers. For longer distances the 10GBASE-SR transceiver can support distances up to 300 meters over OM3 multimode fiber or up to 400 meters over OM4 multimode fiber. The 10GBASE-LR transceivers can support distances up to 10 kilometers on single mode fiber.
To increase the number of available 10 GbE ports, clients can split out four 10 GbE ports for each 40 GbE port using QSFP+ DAC breakout cables for distances up to 5 meters. For distances up to 100 m, the 40GBASE-iSR4 QSFP+ transceivers can be used with OM3 optical MTP-to-LC break-out cables or up to 150 m with OM4 optical MTP-to-LC break-out cables. For longer distances, the 40GBASE-eSR4 transceivers can be used with OM3 optical break-out cables for distances up to 300 m or OM4 optical break-out cables for distances up to 400 m.
- For 40 GbE to 40 GbE connectivity, clients can use the affordable QSFP+ to QSFP+ DAC cables for distances up to 7 meters. For distances up to 100 m, the 40GBASE-SR4/iSR4 QSFP+ transceivers can be used with OM3 multimode fiber or up to 150 m when using OM4 multimode fiber. For distances up to 300 m, the 40GBASE-eSR4 QSFP+ transceiver can be used with OM3 multimode fiber or up to 400 m when using OM4 multimode fiber. For distances up to 10 km, the 40GBASE-LR QSFP+ transceiver can be used with single mode fiber.

Table 3 lists the supported cables and transceivers.

Table 3. Supported transceivers and direct-attach cables

Description	Part number	Feature code	Maximum quantity supported
Serial console cables			
Flex System Management Serial Access Cable Kit	90Y9338	A2RR	1
SFP transceivers - 1 GbE			
Lenovo 1000BASE-T SFP Transceiver (does not support 10/100 Mbps)	00FE333	A5DL	14
Lenovo 1000BASE-SX SFP Transceiver	81Y1622	3269	14
Lenovo 1000BASE-LX SFP Transceiver	90Y9424	A1PN	14
SFP+ transceivers - 10 GbE			
Lenovo 10GBASE-SR SFP+ Transceiver	46C3447	5053	14
Lenovo 10GBASE-LR SFP+ Transceiver	90Y9412	A1PM	14
Optical cables for 1 GbE SX SFP and 10 GbE SR SFP+ transceivers			
Lenovo 1m LC-LC OM3 MMF Cable	00MN502	ASR6	14
Lenovo 3m LC-LC OM3 MMF Cable	00MN505	ASR7	14
Lenovo 5m LC-LC OM3 MMF Cable	00MN508	ASR8	14
Lenovo 10m LC-LC OM3 MMF Cable	00MN511	ASR9	14
Lenovo 15m LC-LC OM3 MMF Cable	00MN514	ASRA	14
Lenovo 25m LC-LC OM3 MMF Cable	00MN517	ASRB	14
Lenovo 30m LC-LC OM3 MMF Cable	00MN520	ASRC	14
SFP+ direct-attach cables - 10 GbE			
Lenovo 1m Passive SFP+ DAC Cable	90Y9427	A1PH	14
Lenovo 1.5m Passive SFP+ DAC Cable	00AY764	A51N	14
Lenovo 2m Passive SFP+ DAC Cable	00AY765	A51P	14
Lenovo 3m Passive SFP+ DAC Cable	90Y9430	A1PJ	14
Lenovo 5m Passive SFP+ DAC Cable	90Y9433	A1PK	14
Lenovo 7m Passive SFP+ DAC Cable	00D6151	A3RH	14
QSFP+ transceiver and cables - 40 GbE			
Lenovo 40GBASE-SR4 QSFP+ Transceiver	49Y7884	A1DR	2
Lenovo 40GBASE-iSR4 QSFP+ Transceiver	00D9865	ASTM	2
Lenovo 40GBASE-eSR4 QSFP+ Transceiver	00FE325	A5U9	2
Lenovo 40GBASE-LR4 QSFP+ Transceiver	00D6222	A3NY	2
Optical cables for 40 GbE QSFP+ SR4/iSR4/eSR4 transceivers			
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable	90Y3519	A1MM	2
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable	90Y3521	A1MN	2
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3519)	00VX003	AT2U	2
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3521)	00VX005	AT2V	2
Optical breakout cables for 40 GbE QSFP+ iSR4/eSR4 transceivers			
Lenovo 1m MTP-4xLC OM3 MMF Breakout Cable	00FM412	A5UA	2
Lenovo 3m MTP-4xLC OM3 MMF Breakout Cable	00FM413	A5UB	2
Lenovo 5m MTP-4xLC OM3 MMF Breakout Cable	00FM414	A5UC	2
QSFP+ breakout cables - 40 GbE to 4x10 GbE			

Description	Part number	Feature code	Maximum quantity supported
Lenovo 1m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7886	A1DL	2
Lenovo 3m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7887	A1DM	2
Lenovo 5m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7888	A1DN	2
QSFP+ direct-attach cables - 40 GbE			
Lenovo 1m Passive QSFP+ DAC Cable	49Y7890	A1DP	2
Lenovo 3m Passive QSFP+ DAC Cable	49Y7891	A1DQ	2
Lenovo 5m Passive QSFP+ DAC Cable	00D5810	A2X8	2
Lenovo 7m Passive QSFP+ DAC Cable	00D5813	A2X9	2

Benefits

The SI4093 interconnect module is considered particularly suited for these clients:

- Clients who want simple 10 GbE network connectivity from the chassis to the upstream network, without the complexity of spanning tree and other advanced Layer 2 and Layer 3 features.
- Clients who want to manage physical compute node connectivity in the chassis by using the existing network management tools.
- Clients looking to build PODs or clusters of up to 252 nodes, but only have to configure once and then can easily scale quickly as needs require.
- Clients who require investment protection for 40 GbE external ports.
- Clients who want to reduce total cost of ownership (TCO) and improve performance, while maintaining high levels of availability and security.
- Clients who want to avoid or minimize oversubscription, which can result in congestion and loss of performance.
- Clients who want to implement a converged infrastructure with NAS, iSCSI, or FCoE. For FCoE implementations, the SI4093 passes through FCoE traffic upstream to other devices, such as the RackSwitch™ G8264CS, Brocade VDX, or Cisco Nexus 5548/5596, where the FC traffic is broken out.

The SI4093 offers the following key features and benefits:

- **Increased performance**
With the growth of virtualization and the evolution of cloud computing, many of today's applications require low latency and high-bandwidth performance. The SI4093 supports submicrosecond latency and up to 1.28 Tbps throughput, while delivering full line rate performance. In addition to supporting 10 GbE ports, the SI4093 can also support 40 GbE external ports, thus enabling forward-thinking clients to connect to their advanced 40 GbE network or as investment protection for the future.

The SI4093 also offers increased security and performance advantage when configured in VLAN-aware mode; it does not force communications upstream into the network, thus reducing latency and generating less network traffic.

- **Pay as you grow investment protection and lower total cost of ownership**
The SI4093 flexible port mapping allows you to buy only the ports that you need, when you need them to lower acquisition and operational costs. The base module configuration includes 24x 10 GbE port licenses that can be assigned to internal connections and 10 GbE or even 40 GbE (by using four 10 GbE licenses per 40 GbE port) external ports. You then have the flexibility of turning on more 10 GbE internal connections and more 10 GbE or 40 GbE external ports when you need them using Lenovo Features on Demand licensing capabilities that provide "pay as you grow" scalability without the need for additional hardware.
- **Cloud ready - Optimized network virtualization with virtual NICs**

With the majority of IT organizations implementing virtualization, there has been an increased need to reduce the cost and complexity of their environments. Lenovo is helping to address these requirements by removing multiple physical I/O ports. Virtual Fabric provides a way for companies to carve up 10 GbE ports into virtual NICs (vNICs) to meet those requirements with Intel processor-based compute nodes.

To help deliver maximum performance per vNIC and to provide higher availability and security with isolation between vNICs, the switch leverages capabilities of its Networking Operating System. For large-scale virtualization, the Flex System solution can support up to 32 vNICs using a pair of CN4054 10Gb Virtual Fabric Adapters in each compute node and four SI4093 modules in the chassis.

The SI4093 switch offers the benefits of next-generation vNIC - Unified Fabric Port (UFP). UFP is an advanced, cost-effective solution that provides a flexible way for clients to allocate, reallocate, and adjust bandwidth to meet their ever-changing data center requirements.

- Cloud ready - VM-aware networking

Delivering advanced virtualization awareness helps simplify management and automates VM mobility by making the network VM aware with VMready, which works with all the major hypervisors. For companies using VMware or KVM, Software Defined Network for Virtual Environment (SDN VE, sold separately) enables network administrators to simplify management by having a consistent virtual and physical networking environment. SDN VE virtual and physical switches use the same configurations, policies, and management tools. Network policies migrate automatically along with virtual machines (VMs) to ensure that security, performance, and access remain intact as VMs move from compute node to compute node.

Support for Edge Virtual Bridging (EVB) based on the IEEE 802.1Qbg standard enables scalable, flexible management of networking configuration and policy requirements per VM and eliminates many of the networking challenges introduced with server virtualization.

- Simplified network infrastructure

The SI4093 simplifies deployment and growth because of its innovative scalable architecture. This architecture helps increase return on investment (ROI) by reducing the qualification cycle, while providing investment protection for additional I/O bandwidth requirements in the future. The extreme flexibility of the interconnect module comes from its ability to turn on additional ports as required, both down to the compute node and for upstream connections (including 40 GbE). Also, as you consider migrating to a converged LAN and SAN, the SI4093 can support the newest protocols, including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE), which can be leveraged in either an iSCSI, Fibre Channel over Ethernet (FCoE), or NAS converged environment.

The SI4093 can be used in the Flex System Interconnect Fabric solution that reduces networking management complexity without compromising performance by lowering the number of devices that need to be managed by 95% (managing one device instead of 20). Interconnect Fabric simplifies POD integration into an upstream network by transparently interconnecting hosts to a data center network and representing the POD as a large compute element isolating the POD's internal connectivity topology and protocols from the rest of the network. Flex System Interconnect Fabric provides simplified and scalable networking infrastructure of up to nine Flex System chassis to build up to a 252 nodes POD or cluster.

The default configuration of the SI4093 requires little or no management for most data center environments, eliminating the need to configure each device or individual ports, thus reducing the number of management points.

Support for Switch Partition (SPAR) allows clients to virtualize the module with partitions that isolate communications for multi-tenancy environments.

- Transparent networking

The SI4093 is a transparent network device, invisible to the upstream network, that eliminates network administration concerns of Spanning Tree Protocol configuration/interoperability, VLAN assignments, and avoidance of possible loops.

By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

Features and specifications

The Flex System Fabric SI4093 System Interconnect Module has the following features and specifications:

- Modes of operations
 - *Transparent (or VLAN-agnostic) mode.*
In VLAN-agnostic mode (default configuration), the SI4093 transparently forwards VLAN tagged frames without filtering on the customer VLAN tag, providing an end host view to the upstream network. The interconnect module provides traffic consolidation in the chassis to minimize TOR port utilization, and it also enables compute node to compute node communication for optimum performance (for example, vMotion). It can be connected to the FCoE transit switch or FCoE gateway (FC Forwarder) device.
 - *Local Domain (or VLAN-aware) mode.*
In VLAN-aware mode (optional configuration), the SI4093 provides additional security for multi-tenant environments by extending client VLAN traffic isolation to the interconnect module and its external ports. VLAN-based access control lists (ACLs) can be configured on the SI4093. When FCoE is used, the SI4093 operates as an FCoE transit switch, and it should be connected to the FCF device.
 - *Flex System Interconnect Fabric mode*
In Flex System Interconnect Fabric mode, the SI4093 module is running optional Interconnect Fabric software image and operates as a leaf switch in the leaf-spine fabric. Flex System Interconnect Fabric integrates the entire point of delivery (POD) into a seamless network fabric for compute node and storage under single IP management, and it attaches to the upstream data center network as a loop-free Layer 2 network fabric with a single Ethernet external connection or aggregation group to each Layer 2 upstream network.
- **Note:** Flexible port mapping is not available in Flex System Interconnect Fabric mode.
- Internal ports
 - Forty-two internal full-duplex 10 Gigabit ports.
 - Two internal full-duplex 1 GbE ports that are connected to the chassis management module.
- External ports
 - Fourteen ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, or 10GBASE-LR) or SFP+ direct-attach copper (DAC) cables. SFP+ modules and DAC cables are not included and must be purchased separately.
 - Two ports for 40 Gb Ethernet QSFP+ transceivers or QSFP+ DAC cables. QSFP+ modules and DAC cables are not included and must be purchased separately.
 - One RS-232 serial port (mini-USB connector) that provides an additional means to configure the interconnect module.
- Scalability and performance
 - 40 Gb Ethernet ports for extreme external bandwidth and performance.
 - External 10 Gb Ethernet ports to leverage 10 GbE upstream infrastructure.
 - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 1.28 Tbps.
 - Media access control (MAC) address learning: automatic update, support for up to 128,000 MAC addresses.
 - Static and LACP (IEEE 802.3ad) link aggregation, up to 220 Gb of total external bandwidth per interconnect module.
 - Support for jumbo frames (up to 9,216 bytes).
- Availability and redundancy
 - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on compute nodes.

- Built-in link redundancy with loop prevention without a need for Spanning Tree protocol.
- VLAN support
 - Up to 32 VLANs supported per interconnect module SPAR partition, with VLAN numbers 1 - 4095. (4095 is used for management module's connection only.)
 - 802.1Q VLAN tagging support on all ports.
 - Private VLANs.
- Security
 - VLAN-based access control lists (ACLs) (VLAN-aware mode).
 - Multiple user IDs and passwords.
 - User access control.
 - Radius, TACACS+, and LDAP authentication and authorization.
 - NIST 800-131A Encryption.
 - Selectable encryption protocol; SHA 256 enabled as default.
- Quality of service (QoS)
 - Support for IEEE 802.1p traffic classification and processing.
- Virtualization
 - Switch Independent Virtual NIC (vNIC2).
 - Ethernet, iSCSI, or FCoE traffic is supported on vNICs.
 - Unified fabric port (UFP)
 - Ethernet or FCoE traffic is supported on UFPs.
 - Supports up to 256 VLAN for the virtual ports.
 - Integration with L2 failover.
 - 802.1Qbg Edge Virtual Bridging (EVB) is an emerging IEEE standard for allowing networks to become virtual machine (VM)-aware.
 - Virtual Ethernet Bridging (VEB) and Virtual Ethernet Port Aggregator (VEPA) are mechanisms for switching between VMs on the same hypervisor.
 - Edge Control Protocol (ECP) is a transport protocol that operates between two peers over an IEEE 802 LAN providing reliable, in-order delivery of upper layer protocol data units.
 - Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) allows centralized configuration of network policies that will persist with the VM, independent of its location.
 - EVB Type-Length-Value (TLV) is used to discover and configure VEPA, ECP, and VDP.
 - VMready
 - Switch partitioning (SPAR)
 - SPAR forms separate virtual switching contexts by segmenting the data plane of the module. Data plane traffic is not shared between SPARs on the same module.
 - SPAR operates as a Layer 2 broadcast network. Hosts on the same VLAN attached to a SPAR can communicate with each other and with the upstream switch. Hosts on the same VLAN but attached to different SPARs communicate through the upstream switch.
 - SPAR is implemented as a dedicated VLAN with a set of internal compute node ports and a single external port or link aggregation (LAG). Multiple external ports or LAGs are not allowed in SPAR. A port can be a member of only one SPAR.
- Converged Enhanced Ethernet
 - Priority-Based Flow Control (PFC) (IEEE 802.1Qbb) extends 802.3x standard flow control to allow the module to pause traffic based on the 802.1p priority value in each packet's VLAN tag.
 - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz) provides a method for allocating link bandwidth based on the 802.1p priority value in each packet's VLAN tag.
 - Data Center Bridging Capability Exchange Protocol (DCBX) (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities.
- Fibre Channel over Ethernet (FCoE)
 - FC-BB5 FCoE specification compliant.
 - FCoE transit switch operations.
 - FCoE Initialization Protocol (FIP) support.
- Manageability

- IPv4 and IPv6 host management.
- Simple Network Management Protocol (SNMP V1, V2, and V3).
- Industry standard command-line interface (IS-CLI) through Telnet, SSH, and serial port.
- Secure FTP (sFTP).
- Service Location Protocol (SLP).
- Firmware image update (TFTP and FTP/sFTP).
- Network Time Protocol (NTP) for clock synchronization.
- Switch Center support.
- Monitoring
 - LEDs for external port status and module status indication.
 - Change tracking and remote logging with syslog feature.
 - POST diagnostic tests.

Standards supported

The SI4093 supports the following standards:

- IEEE 802.1AB Data Center Bridging Capability Exchange Protocol (DCBX)
- IEEE 802.1p Class of Service (CoS) prioritization
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1Qbb Priority-Based Flow Control (PFC)
- IEEE 802.1Qaz Enhanced Transmission Selection (ETS)
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3ab 1000BASE-T copper twisted pair Gigabit Ethernet
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3ae 10GBASE-SR short range fiber optics 10 Gb Ethernet
- IEEE 802.3ae 10GBASE-LR long range fiber optics 10 Gb Ethernet
- IEEE 802.3ap 10GBASE-KR backplane 10 Gb Ethernet
- IEEE 802.3ba 40GBASE-SR4 short range fiber optics 40 Gb Ethernet
- IEEE 802.3ba 40GBASE-CR4 copper 40 Gb Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3z 1000BASE-SX short range fiber optics Gigabit Ethernet
- IEEE 802.3z 1000BASE-LX long range fiber optics Gigabit Ethernet
- SFF-8431 10GSFP+Cu SFP+ Direct Attach Cable

Chassis and adapters

The interconnect modules are installed in I/O module bays in the rear of the Flex System chassis, as shown in Figure 2. Switches are normally installed in pairs because ports on the I/O adapter cards installed in the compute nodes are routed to two I/O bays for redundancy and performance.




Figure 2. Location of the I/O bays in the Flex System chassis

The SI4093 module can be installed in bays 1, 2, 3, and 4 of the Flex System chassis. A supported adapter card must be installed in the corresponding slot of the compute node. Each adapter can use up to four lanes to connect to the respective I/O module bay. The SI4093 is able to use up to three of the four lanes.

Prior to Networking OS 7.8, with four-port adapters, an optional Upgrade 1 (95Y3318) was required for the SI4093s to allow communications on all four ports. With eight-port adapters, both optional Upgrade 1 (95Y3318) and Upgrade 2 (95Y3320) were required for the module to allow communications on six adapter ports, and two remaining ports are not used. With Networking OS 7.8 or later, there is no need to buy additional module upgrades for 4-port and 8-port adapters if the total number of port licenses on the SI4093 does not exceed the number of external (upstream network ports) and internal (compute node network ports) connections used.

In compute nodes that have an integrated dual-port 10 GbE network interface controller (NIC), NIC ports are routed to bays 1 and 2 with a specialized periscope connector, and the adapter card is not required. However, when needed, the periscope connector can be replaced with the adapter card. In this case, integrated NIC will be disabled.

Table 4 shows compatibility information for the SI4093 and Flex System chassis.

Table 4. Flex System chassis compatibility

Description	Part number	Enterprise Chassis with CMM	Enterprise Chassis with CMM2	Carrier-grade Chassis with CMM2
Flex System Fabric SI4093 System Interconnect Module	95Y3313	Yes	Yes	No

The midplane connections between the adapters installed in the compute nodes to the switch bays in the chassis are shown diagrammatically in the following figure. The figure shows both half-wide compute nodes, such as the x240 with two adapters, and full-wide compute nodes, such as the x440 with four adapters.




Figure 3. Logical layout of the interconnects between I/O adapters and I/O modules

Table 5 shows the connections between adapters installed in the compute nodes to the I/O bays in the chassis.

Table 5. Adapter to I/O bay correspondence

I/O adapter slot in the compute node	Port on the adapter	Corresponding I/O module bay in the chassis			
		Bay 1	Bay 2	Bay 3	Bay 4
Slot 1	Port 1	Yes			
	Port 2		Yes		
	Port 3	Yes			
	Port 4		Yes		
	Port 5	Yes			
	Port 6		Yes		
	Port 7*				
	Port 8*				
Slot 2	Port 1			Yes	
	Port 2				Yes
	Port 3			Yes	
	Port 4				Yes
	Port 5			Yes	
	Port 6				Yes
	Port 7*				
	Port 8*				
Slot 3 (full-wide compute nodes only)	Port 1	Yes			
	Port 2		Yes		
	Port 3	Yes			
	Port 4		Yes		
	Port 5	Yes			
	Port 6		Yes		
	Port 7*				
	Port 8*				
Slot 4 (full-wide compute nodes only)	Port 1			Yes	
	Port 2				Yes
	Port 3			Yes	
	Port 4				Yes
	Port 5			Yes	
	Port 6				Yes
	Port 7*				
	Port 8*				

* Ports 7 and 8 are routed to I/O bays 1 and 2 (Slot 1 and Slot 3) or 3 and 4 (Slot 2 and Slot 4), but these ports cannot be used with the SI4093.

The following table lists the adapters that are supported by the I/O module.

Table 6. Network adapters

Description	Part number	Feature code
50 Gb Ethernet		
ThinkSystem QLogic QL45212 Flex 50Gb 2-Port Ethernet Adapter	7XC7A05843	B2VT
ThinkSystem QLogic QL45262 Flex 50Gb 2-Port Ethernet Adapter with iSCSI/FCoE	7XC7A05845	B2VV
25 Gb Ethernet		
ThinkSystem QLogic QL45214 Flex 25Gb 4-Port Ethernet Adapter	7XC7A05844	B2VU
10 Gb Ethernet		
Embedded 10Gb Virtual Fabric Adapter (2-port)†	None	None
Flex System CN4022 2-port 10Gb Converged Adapter	88Y5920	A4K3
Flex System CN4052 2-port 10Gb Virtual Fabric Adapter	00JY800*	A5RP
Flex System CN4052S 2-port 10Gb Virtual Fabric Adapter	00AG540	ATBT
Flex System CN4052S 2-port 10Gb Virtual Fabric Adapter Advanced	01CV780	AU7X
Flex System CN4054 10Gb Virtual Fabric Adapter (4-port)	90Y3554*	A1R1
Flex System CN4054R 10Gb Virtual Fabric Adapter (4-port)	00Y3306*	A4K2
Flex System CN4054S 4-port 10Gb Virtual Fabric Adapter	00AG590	ATBS
Flex System CN4054S 4-port 10Gb Virtual Fabric Adapter Advanced	01CV790	AU7Y
Flex System CN4058S 8-port 10Gb Virtual Fabric Adapter	94Y5160	A4R6
Flex System EN4172 2-port 10Gb Ethernet Adapter	00AG530	A5RN
1 Gb Ethernet		
Embedded 1 Gb Ethernet controller (2-port)**	None	None
Flex System EN2024 4-port 1Gb Ethernet Adapter	49Y7900	A10Y

* Withdrawn from marketing

† The Embedded 10Gb Virtual Fabric Adapter is built into selected compute nodes.

** The Embedded 1 Gb Ethernet controller is built into selected compute nodes.

The adapters are installed in slots in each compute node. Figure 4 shows the locations of the slots in the x240 Compute Node. The positions of the adapters in the other supported compute nodes are similar.




Figure 4. Location of the I/O adapter slots in the Flex System x240 Compute Node

Connectors and LEDs

Figure 5 shows the front panel of the Flex System Fabric SI4093 System Interconnect Module.

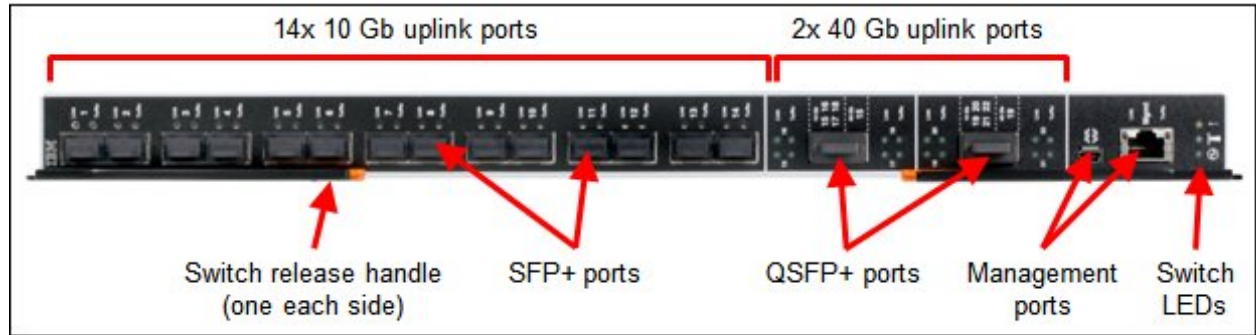


Figure 5. Front panel of the Flex System Fabric SI4093 System Interconnect Module

The front panel contains the following components:

- LEDs that display the status of the interconnect module and the network:
 - The OK LED indicates that the interconnect module passed the power-on self-test (POST) with no critical faults and is operational.
 - Identify: This blue LED can be used to identify the module physically by illuminating it through the management software.
 - The error LED (module error) indicates that the module failed the POST or detected an operational fault.
- One mini-USB RS-232 console port that provides an additional means to configure the interconnect module. This mini-USB-style connector enables connection of a special serial cable. (The cable is optional and it is not included with the interconnect module. For more information, see the "Part number information" section.)
- Fourteen external SFP+ ports for 1 GbE or 10 GbE connections to external Ethernet devices.
- Two external QSFP+ port connectors to attach QSFP+ transceivers or cables for a single 40 GbE connectivity or for splitting of a single port into 4x 10 GbE connections.
- An Ethernet link OK LED and an Ethernet Tx/Rx LED for each external port on the interconnect module.

Network cabling requirements

The network cables that can be used with the SI4093 module are shown in Table 7.

Table 7. SI4093 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
40GBASE-SR4 QSFP+ Transceiver (49Y7884)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables supplied by Lenovo (see Table 3); support for up to 100 m with OM3 multimode fiber or up to 150 m with OM4 multimode fiber	MTP
40GBASE-iSR4 QSFP+ Transceiver (00D9865)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables or MTP-4xLC breakout cables up to 5 m supplied by Lenovo (see Table 3); support for up to 100 m with OM3 multimode fiber or up to 150 m with OM4 multimode fiber	MTP
40GBASE-eSR4 QSFP+ Transceiver (00FE325)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables or MTP-4xLC breakout cables up to 5 m supplied by Lenovo (see Table 3); support for up to 300 m with OM3 multimode fiber or up to 400 m with OM4 multimode fiber	MTP
40GBASE-LR4 QSFP+ Transceiver (00D6222)	40GBASE-LR4	1310 nm single-mode fiber cable up to 10 km	LC
Direct attach cable	40GBASE-CR4	QSFP+ to QSFP+ DAC cables up to 7 m; QSFP+ to 4x SFP+ DAC break-out cables up to 5 m for 4x 10 GbE SFP+ connections out of a 40 GbE port (see Table 3)	QSFP+
10 Gb Ethernet			
10GBASE-SR SFP+ Transceiver (46C3447)	10GBASE-SR	Up to 30 m with fiber optic cables supplied by Lenovo (see Table 3); 850 nm OM3 multimode fiber cable up to 300 m or up to 400 m with OM4 multimode fiber	LC
10GBASE-LR SFP+ Transceiver (90Y9412)	10GBASE-LR	1310 nm single-mode fiber cable up to 10 km	LC
Direct attach cable	10GSFP+Cu	SFP+ DAC cables up to 7 m (see Table 3)	SFP+
1 Gb Ethernet			
1000BASE-T SFP Transceiver (00FE333)	1000BASE-T	UTP Category 5, 5E, and 6 up to 100 meters	RJ-45
1000BASE-SX SFP Transceiver (81Y1622)	1000BASE-SX	Up to 30 m with fiber optic cables supplied by Lenovo (see Table 3); 850 nm multimode fiber cable up to 550 m (50 μ) or up to 220 m (62.5 μ)	LC
1000BASE-LX SFP Transceiver (90Y9424)	1000BASE-LX	1310 nm single-mode fiber cable up to 10 km	LC
Management ports			
1 GbE management port	1000BASE-T	UTP Category 5, 5E, and 6 up to 100 meters	RJ-45
External RS-232 management port	RS-232	DB-9-to-mini-USB or RJ-45-to-mini-USB console cable (comes with optional Management Serial Access Cable, 90Y9338)	Mini-USB

Warranty

The SI4093 carries a 1-year, customer-replaceable unit (CRU) limited warranty. When installed in a chassis, these I/O modules assume your system's base warranty and any warranty service upgrade.

Physical specifications

Here are the approximate dimensions and weight of the SI4093:

- Height: 30 mm (1.2 in.)
- Width: 401 mm (15.8 in.)
- Depth: 317 mm (12.5 in.)
- Weight: 3.7 kg (8.1 lb)

Shipping dimensions and weight (approximate):

- Height: 114 mm (4.5 in.)
- Width: 508 mm (20.0 in.)
- Depth: 432 mm (17.0 in.)
- Weight: 4.1 kg (9.1 lb)

Agency approvals

The SI4093 conforms to the following regulations:

- United States FCC 47 CFR Part 15, Subpart B, ANSI C63.4 (2003), Class A
- IEC/EN 60950-1, Second Edition
- Canada ICES-003, issue 4, Class A
- Japan VCCI, Class A
- Australia/New Zealand AS/NZS CISPR 22:2006, Class A
- Taiwan BSMI CNS13438, Class A
- CE Mark (EN55022 Class A, EN55024, EN61000-3-2, EN61000-3-3)
- CISPR 22, Class A
- China GB 9254-1998
- Turkey Communiqué 2004/9; Communiqué 2004/22
- Saudi Arabia EMC.CVG, 28 October 2002

Typical configurations

The most common SI4093 connectivity topology, which can be used with both Lenovo and non-Lenovo upstream network devices, is shown in Figure 6.

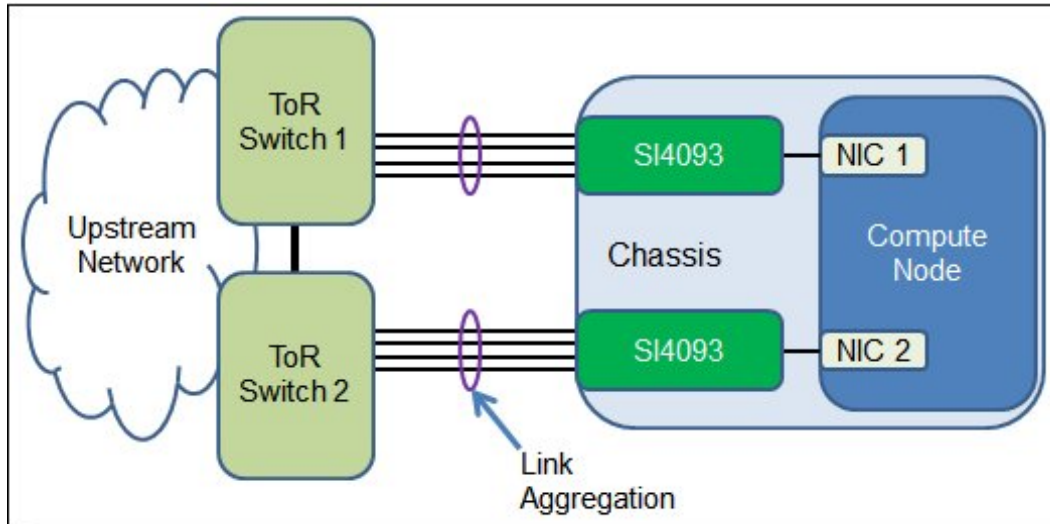


Figure 6. SI4093 connectivity topology - Link Aggregation

In this loop-free redundant topology, each SI4093 is physically connected to a separate Top-of-Rack (ToR) switch with static or LACP aggregated links.

When the SI4093 is used with the RackSwitch switches, Virtual Link Aggregation Groups (vLAGs) can be used for load balancing and redundancy purposes. The virtual link aggregation topology is shown in Figure 7.

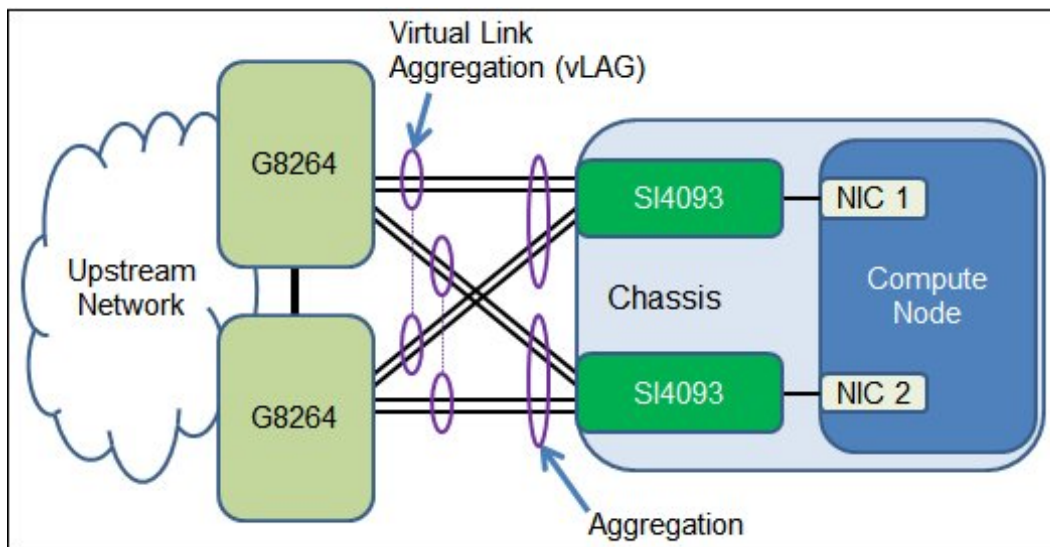


Figure 7. SI4093 connectivity topology - Virtual Link Aggregation

In this loop-free topology, aggregation is split between two physical switches, which appear as a single logical switch, and each SI4093 is connected to both ToR switches through static or LACP aggregated links.

Dual isolated SAN fabrics: If you plan to use FCoE and follow a dual isolated SAN fabric design approach (also known as *SAN air gaps*), consider the SI4093 connectivity topology shown in Figure 6 (Link Aggregation).

The following usage scenarios are considered:

- SI4093 in the traditional 10 Gb Ethernet network
- SI4093 in the converged FCoE network
- SI4093 in the Flex System Interconnect Fabric solution

SI4093 in the traditional 10 Gb Ethernet network

In the traditional 10 GbE network, the SI4093 can be used together with pNIC or the switch-independent vNIC capabilities of the 10 Gb Virtual Fabric Adapters (VFAs) that are installed in each compute node. In the case of vNIC, each physical port on the adapter is split into four virtual NICs (vNICs) with a certain amount of bandwidth that is assigned. vNIC bandwidth allocation and metering is performed by the VFA, and a unidirectional virtual channel of an assigned bandwidth is established between the I/O module and the VFA for each vNIC. Up to 32 vNICs can be configured on a half-wide compute node.

The SI4093 interconnect modules is connected to the ToR aggregator switches, such as the following ones:

- RackSwitch G8264 through the 10 GbE SFP+ ports or QSFP+ break out cables
- RackSwitch G8316/G8332 through the 40 GbE QSFP+ ports

Figure 8 illustrates this scenario.




Figure 8. SI4093 in the 10 GbE network

The solution components that are used in the scenario that is shown in Figure 8 are listed in Table 8.

Table 8. Components that are used in 10 GbE solution with the SI4093 (Figure 8)

Diagram reference	Description	Part number	Quantity
1	Flex System Virtual Fabric solution		
	Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per compute node
	Flex System Fabric SI4093 System Interconnect Module	95Y3313	2 per chassis
	Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)*	95Y3318	1 per SI4093
2	RackSwitch G8264, G8316, or G8332		

* The Upgrade 1 might not be needed with flexible port mapping if the total number of the internal and external ports used on the SI4093 is less or equal to 24.

Note: You also need SFP+/QSFP+ modules and optical cables or SFP+/QSFP+ DAC cables (not shown in Table 8; see Table 3 for details) for the external 10 Gb Ethernet connectivity.

SI4093 in the converged FCoE network

SI4093 supports Data Center Bridging (DCB), and it can transport FCoE frames. These interconnect modules provide an inexpensive solution for transporting encapsulated FCoE packets to the Fibre Channel Forwarder (FCF), which is functioning as both an aggregation switch and an FCoE gateway. Vendor-specific examples of this scenario are shown in Figure 9, Figure 10, and Figure 11. The solution components that are used in the scenarios that are shown in Figure 9, Figure 10, and Figure 11 are listed in Table 9, Table 10, and Table 11, respectively.




Figure 9. SI4093 in the FCoE network with the RackSwitch G8264CS as an FCF




Figure 10. SI4093 in the FCoE network with the Brocade VDX 6730 as an FCF

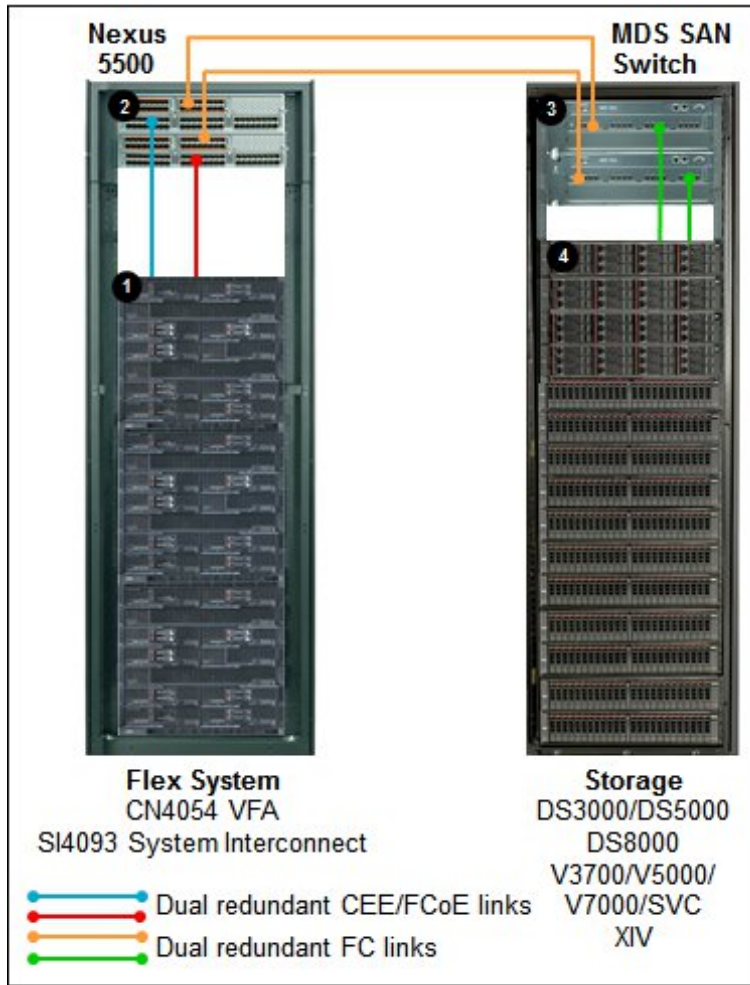


Figure 11. SI4093 in the FCoE network with the Cisco Nexus 5548/5596 as an FCF

Table 9. SI4093 with the G8264CS as an FCF (Figure 9)

Diagram reference	Description	Part number	Quantity
1	Flex System FCoE solution		
	Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per compute node
	Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	Flex System Fabric SI4093 System Interconnect Module	95Y3313	2 per chassis
	Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)*	95Y3318	1 per SI4093
2	RackSwitch G8264CS		
3	Brocade or Cisco MDS SAN fabric		
4	Storage systems		
	IBM DS3000 / DS5000		
	IBM DS8000		
	IBM Storwize V3700 / V5000 / V7000 / SAN Volume Controller		
	IBM XIV		

* The Upgrade 1 might not be needed with flexible port mapping if the total number of the internal and external ports used on the SI4093 is less or equal to 24.

Table 10. SI4093 with the Brocade VDX 6730 as an FCF (Figure 10)

Diagram reference	Description	Part number	Quantity
1	Flex System FCoE solution		
	Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per compute node
	Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	Flex System Fabric SI4093 System Interconnect Module	95Y3313	2 per chassis
	Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)*	95Y3318	1 per SI4093
2	Brocade VDX 6730 Converged Switch		
3	Brocade SAN fabric		
4	Storage systems		
	DS3000 / DS5000		
	DS8000®		
	Storwize V3700 / V5000 / V7000 / SAN Volume Controller		
	XIV		

* The Upgrade 1 might not be needed with flexible port mapping if the total number of the internal and external ports used on the SI4093 is less or equal to 24.

Table 11. SI4093 with the Cisco Nexus 5548/5596 as an FCF (Figure 11)

Diagram reference	Description	Part number	Quantity
1	Flex System FCoE solution		
	Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per compute node
	Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	Flex System Fabric SI4093 System Interconnect Module	95Y3313	2 per chassis
	Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)*	95Y3318	1 per SI4093
2	Cisco Nexus 5548/5596 Switch		
3	Cisco MDS SAN fabric		
4	Storage systems		
	DS3500 / DS5000		
	DS8000®		
	Storwize V3700 / V5000 / V7000 / SAN Volume Controller		
	XIV		

* The Upgrade 1 might not be needed with flexible port mapping if the total number of the internal and external ports used on the SI4093 is less or equal to 24.

Note: You also need SFP+ modules and optical cables or SFP+ DAC cables (not shown in Table 9, Table 10, and Table 11; see Table 3 for details) for the external 10 Gb Ethernet connectivity.

Lenovo provides extensive FCoE testing to deliver network interoperability. For a full listing of supported FCoE and iSCSI configurations, see the System Storage Interoperation Center (SSIC) website at:

<http://ibm.com/systems/support/storage/ssic>

SI4093 in the Flex System Interconnect Fabric solution

A typical Flex System Interconnect Fabric configuration uses four 10 GbE ports as Fabric Ports from each SI4093, two ports to each of the aggregation switches. The maximum of nine chassis, with a total of eighteen SI4093 modules use 36 ports on each of the G8264CS aggregation switches. It is possible to use additional ports from the SI4093s to the aggregation switches if there are ports available, up to a total of eight.

The Interconnect Fabric can be VLAN-aware or VLAN-agnostic depending on specific client requirements:

- In VLAN-aware mode, client VLAN isolation is extended to the fabric by filtering and forwarding VLAN tagged frames based on the client VLAN tag, and client VLANs from the upstream network are configured within the Interconnect Fabric.
- In VLAN-agnostic mode, the Interconnect Fabric transparently forwards VLAN tagged frames without filtering on the client VLAN tag, providing an end host view to the upstream network. This is achieved by the use of a Q-in-Q type operation to hide user VLANs from the switch fabric in the POD so that the Interconnect Fabric acts as more of a port aggregator and is user VLAN-independent.

As an example, Figure 12 shows a 3-chassis Interconnect Fabric POD consisting of:

- 6x SI4093 modules
- 2x G8264CS switches with required SFP+ DAC cables (four 10 GbE SFP+ DAC cables per one SI4093 plus two 10 GbE SFP+ DAC cables to connect G8264CS switches to each other)
- Required rack and PDU infrastructure
- Optional Switch Center management application




Figure 12. SI4093 in the Flex System Interconnect Fabric solution

The solution components that are used in the scenario that is shown in Figure 12 is listed in Table 12.

Table 12. SI4093 in the Flex System Interconnect Fabric solution (Figure 12)

Diagram reference	Description	Part number	Quantity
1	Rack and PDU infrastructure		
	42U 1200mm Deep Dynamic Rack	93604PX	1
	0U 12 C19/12 C13 32A 3 Phase PDU	46M4143	2
	1U Quick Install Filler Panel Kit	25R5559	2
2	Top of Rack switches		
	RackSwitch G8264CS (Rear-to-Front)	7309DRX	2
	3m Passive DAC SFP+ Cable	90Y9430	16
	1m Passive DAC SFP+ Cable	90Y9427	10
3	Flex System Enterprise Chassis with SI4093 modules		
	Flex System Enterprise Chassis with 2x2500W PSU	8721A1G	3
	Flex System Enterprise Chassis 2500W Power Module	43W9049	12
	Flex System Fabric SI4093 System Interconnect Module	95Y3313	6
	Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)*	95Y3318	6
	Flex System Chassis Management Module	68Y7030	3
	Flex System Enterprise Chassis 80mm Fan Module Pair	43W9078	6
	Management application (optional)		
	Switch Center, per install with 1 year software subscription and support for 20 switches	00AE226	1

* Upgrade 1 is required for x222 compute nodes or compute nodes with 4-port network adapters installed. Upgrade 1 is not required if 2-port LOM on the compute node other than x222 is used for network connectivity.

Note: Cables or SFP+ modules for the upstream network connectivity are not included.

Related publications and links

For more information, see the following Flex System Fabric SI4093 System Interconnect Module product publications, which are available from the Flex System Information Center at <http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>:

- *Installation Guide*
- *Application Guide*
- *Command Reference*

Here are other useful references:

- US Announcement Letter for the SI4093 Interconnect Module
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS113-133>
- *Flex System Enterprise Chassis* Product Guide
<http://lenovopress/tips0865>
- *Flex System Products and Technology*, SG24-8255
<http://lenovopress/sg248255>
- Flex System Interoperability Guide
<http://lenovopress.com/fsig>

Related product families

Product families related to this document are the following:

- [10 Gb Embedded Connectivity](#)
- [Blade Networking Modules](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.

Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2023. All rights reserved.

This document, TIPS1045, was created or updated on May 7, 2018.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/TIPS1045>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS1045>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
Flex System
RackSwitch
ThinkSystem®
VMready®

The following terms are trademarks of other companies:

Intel® is a trademark of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.