

Flex System Fabric CN4093 10Gb Converged Scalable Switch

Product Guide (withdrawn product)

The Flex System® Fabric CN4093 10Gb Converged Scalable Switch provides unmatched scalability, port flexibility, performance, convergence, and network virtualization, while also delivering innovations to help address a number of networking concerns today and providing capabilities that will help you prepare for the future. The switch offers full Layer 2/3 switching as well as FCoE Full Fabric and Fibre Channel NPV Gateway operations to deliver a truly converged integrated solution, and it is designed to install within the I/O module bays of the Flex System Enterprise Chassis. The switch can help clients migrate to a 10 GbE or 40 GbE converged Ethernet infrastructure and offers virtualization features like Virtual Fabric and VMready®.

Figure 1 shows the Flex System Fabric CN4093 10Gb Converged Scalable Switch.



Figure 1. Flex System Fabric CN4093 10Gb Converged Scalable Switch

Did you know?

The CN4093 offers up to 12 external Omni Ports, which provide extreme flexibility with the choice of SFP+ based 10 Gb Ethernet connectivity or 4/8 Gb Fibre Channel connectivity, depending on the SFP+ module used.

The base switch configuration comes standard with 22x 10 GbE port licenses that can be assigned to internal connections or external SFP+, Omni 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+ or Omni 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 Lenovo Features on Demand upgrade licensing capabilities that provide “pay as you grow” scalability without the need to buy additional hardware.

Part number information

The CN4093 switch is initially licensed for 22x 10 GbE ports. Further ports can be enabled with Upgrade 1 and Upgrade 2 license options. Upgrade 1 and Upgrade 2 can be applied on the switch independently from each other or in combination for full feature capability. Table 1 shows the part numbers for ordering the switches and the upgrades.

Table 1. Part numbers and feature codes for ordering

Description	Part number	Feature code
Switch module		
Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	A3HH
Features on Demand upgrades		
Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845	A3HL
Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 2)	00D5847	A3HM

The part number for the switch includes the following items:

- One Flex System Fabric CN4093 10Gb Converged Scalable Switch
- Documentation package

Note: Neither QSFP+ nor SFP+ transceivers or cables are included with the switch. They must be ordered separately (See Table 3).

The switch 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 switch locally for configuration tasks and firmware updates.

The part numbers for the upgrades, 00D5845 and 00D5847, include the following items:

- Features on Demand Activation Flyer
- Upgrade authorization letter

The base switch and upgrades are as follows:

- 00D5823 is the part number for the base switch, and it comes with 14 internal 10 GbE ports enabled (one to each node bay), two external 10 GbE SFP+ ports enabled, and six Omni Ports enabled to connect to either Ethernet or Fibre Channel networking infrastructure, depending on the SFP+ transceiver or DAC cable used.
- 00D5845 (Upgrade 1) can be applied on the base switch when you need more external bandwidth with two 40 GbE QSFP+ ports that can be converted into 4x 10 GbE SFP+ links each with the optional break-out cables. This upgrade also enables 14 additional internal ports, for a total of 28 internal ports, to provide more bandwidth to the compute nodes leveraging 4-port expansion cards. This takes full advantage of four-port adapter cards installed in each compute node and requires the base switch.
- 00D5847 (Upgrade 2) can be applied on the base switch when you need more external Omni Ports on the switch or if you want additional internal bandwidth to the node bays. The upgrade will enable the remaining six external Omni Ports, plus 14 additional internal 10 GbE ports, for a total of 28 internal ports, to provide more bandwidth to the compute nodes leveraging four-port expansion cards. This takes full advantage of four-port adapter cards installed in each compute node and requires the base switch.
- Both 00D5845 (Upgrade 1) and 00D5847 (Upgrade 2) can be applied on the switch at the same time to allow you to use 42 internal 10 GbE ports leveraging six ports on an eight-port expansion card, and to utilize all external ports on the switch.

Flexible port mapping: With Networking OS version 7.8 or later clients have more flexibility in assigning ports that they have licensed on the CN4093 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 and Omni 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 or Upgrade 2.

Note: Flexible port mapping is not available in Stacking mode.

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

- 00D5823 is the part number for the base switch, and it provides 22x 10 GbE port licenses that can enable any combination of internal and external 10 GbE ports and Omni Ports and external 40 GbE ports (with the use of four 10 GbE port licenses per one 40 GbE port).
- 00D5845 (Upgrade 1) upgrades the base switch 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 44x 10 GbE port licenses. Any combination of internal and external 10 GbE ports and Omni 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 switch.
- 00D5847 (Upgrade 2) upgrades the base switch by activation of 14 internal 10 GbE ports and six external Omni Ports which is equivalent to adding 20 more 10 GbE port licenses for a total of 42x 10 GbE port licenses. Any combination of internal and external 10 GbE ports and Omni 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 switch.
- Both 00D5845 (Upgrade 1) and 00D5847 (Upgrade 2) simply activate all the ports on the CN4093 which is 42 internal 10 GbE ports, two external SFP+ ports, 12 external Omni 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 CN4093 are enabled.

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

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

Supported port combinations	Quantity required		
	Base switch, 00D5823	Upgrade 1, 00D5845	Upgrade 2, 00D5847
<ul style="list-style-type: none"> • 14x internal 10 GbE ports • 2x external 10 GbE SFP+ ports • 6x external SFP+ Omni Ports 	1	0	0
<ul style="list-style-type: none"> • 28x internal 10 GbE ports • 2x external 10 GbE SFP+ ports • 6x external SFP+ Omni Ports • 2x external 40 GbE QSFP+ ports 	1	1	0
<ul style="list-style-type: none"> • 28x internal 10 GbE ports • 2x external 10 GbE SFP+ ports • 12x external SFP+ Omni Ports 	1	0	1
<ul style="list-style-type: none"> • 42x internal 10 GbE ports† • 2x external 10 GbE SFP+ ports • 12x external SFP+ Omni Ports • 2x external 40 GbE QSFP+ 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	Quantity required		
	Base switch, 00D5823	Upgrade 1, 00D5845	Upgrade 2, 00D5847
<ul style="list-style-type: none"> • 22x 10 GbE ports (internal and external SFP+ and Omni Ports) or • 18x 10 GbE ports (internal and external SFP+ and Omni Ports) • 1x external 40 GbE QSFP+ ports or • 14x 10 GbE ports (internal and external SFP+ and Omni Ports) • 2x external 40 GbE QSFP+ ports 	1	0	0
<ul style="list-style-type: none"> • 44x 10 GbE ports (internal and external SFP+ and Omni Ports) or • 40x 10 GbE ports (internal and external SFP+ and Omni Ports) • 1x external 40 GbE QSFP+ ports or • 36x 10 GbE ports (internal and external SFP+ and Omni Ports) • 2x external 40 GbE QSFP+ ports 	1	1	0
<ul style="list-style-type: none"> • 42x 10 GbE ports (internal and external SFP+ and Omni Ports) or • 38x 10 GbE ports (internal and external SFP+ and Omni Ports) • 1x external 40 GbE QSFP+ ports or • 34x 10 GbE ports (internal and external SFP+ and Omni Ports) • 2x external 40 GbE QSFP+ ports 	1	0	1

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

Supported transceivers and cables

With the flexibility of the CN4093 switch, 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 (supported on external SFP+ ports), you can use SFP+ direct attach copper (DAC) cables for in-rack cabling and distances up to 7 m (up to 5 m for Omni Ports). 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.
- For 4 Gb or 8 Gb FC links (supported on Omni Ports only), you can use 8 Gb FC SFP+ SW optical transceivers plus LC fiber optics cables for distances up to 150 m with 50 μ multi-mode fiber or up to 21 m with 62.5 μ multi-mode fiber. These transceivers can operate at 4 Gb or 8 Gb speeds.

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	2
Lenovo 1000BASE-SX SFP Transceiver	81Y1622	3269	2
Lenovo 1000BASE-LX SFP Transceiver	90Y9424	A1PN	2
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, 10 GbE SR SFP+, and 8 Gb FC SW 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	2*
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			

Description	Part number	Feature code	Maximum quantity supported
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			
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
SFP+ transceivers - 8 Gb FC (supported in Omni Ports)			
8Gb SFP+ SW Optical Transceiver (supports 4/8 Gbps)	44X1964	5075	12

* Supported in SFP+ ports only, not supported in Omni Ports.

Benefits

The Flex System Fabric CN4093 10Gb Converged Scalable Switch is considered particularly suited for these clients:

- Clients who want to implement a converged infrastructure with FCoE where the CN4093 acts as a Full Fabric FC/FCoE switch for the end-to-end FCoE configurations or as an integrated Fibre Channel Forwarder (FCF) NPV Gateway breaking out FC traffic within the chassis for the native Fibre Channel SAN connectivity.
- Clients who are implementing a virtualized environment.
- Clients who require investment protection for 40 GbE external ports.
- Clients who want to reduce 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.

The switch offers the following key capabilities and benefits:

- **Convergence and lower acquisition and operational costs**
One of the key trends that is driving the transformation of the data center is converging to a simplified networking infrastructure — collapsing Ethernet and Fibre Channel at the server and the edge of the network while maintaining connectivity upstream to existing LANs and SANs. The CN4093 Converged Switch supports multiple protocols, including Ethernet, Fibre Channel, iSCSI, and FCoE; Omni Ports give clients the flexibility to choose between 10 Gb Ethernet external connections to the top-of-rack switch and 4/8 Gb Fibre Channel for flexible and scalable access to FC storage.

- Increased performance**
 With the growth of virtualization and the evolution of cloud, many of today's applications require low latency and high bandwidth performance. The CN4093 is the embedded 10 GbE switch for a Flex System chassis to support sub-microsecond latency and up to 1.28 Tbps, while also delivering full line rate performance on Ethernet ports, making it ideal for managing dynamic workloads across the network. In addition, this switch provides a rich Layer 2 and Layer 3 feature set that is ideal for many of today's data centers. Furthermore, it offers industry-leading external bandwidth by being the integrated switch to support 40 GbE external ports.
- Pay as you grow investment protection and lower total cost of ownership**
 The CN4093's flexible port mapping allows you to buy only the ports that you need, when you need them. The base switch configuration includes 22x 10 GbE port licenses that can be assigned to internal connections and external SFP+, Omni or 40 GbE QSFP+ ports (by using four 10 GbE licenses per one 40 GbE port). You then have the flexibility of turning on more internal 10 GbE connections and more external ports when you need them using 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 by using a pair of CN4054 10Gb Virtual Fabric Adapters in each compute node and four CN4093 10Gb Converged Scalable Switches in the chassis.

 The CN4093 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, SDN for Virtual Environments (sold separately) enables network administrators to simplify management by having a consistent virtual and physical networking environment. With 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.
- Simplify network infrastructure**
 The CN4093 10Gb Converged Scalable Switch simplifies deployment and growth by using its innovative scalable architecture. This architecture helps increase return on investment by reducing the qualification cycle, while providing investment protection for additional I/O bandwidth requirements in the future. The extreme flexibility of the switch comes from the 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 CN4093 supports Omni Ports for Ethernet or FC connectivity, and it can operate as an integrated FCF, which can be leveraged in an FCoE converged environment.

CN4093 hybrid stacking capabilities simplify management for clients by stacking up to eight switches (two CN4093 and from two to six EN4093/EN4093R switches) that share one IP address and one management interface. Support for Switch Partition (SPAR) allows clients to virtualize the switch with partitions that isolate communications for multi-tenancy environments.

- **Transparent networking capability**
With a simple configuration change to "easy connect" mode, the CN4093 becomes a transparent network device, invisible to the core, eliminating 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.

- **Integrate network management**
A key challenge is the management of a discrete network environment. The CN4093 supports a command-line interface (CLI) for integration into existing scripting and automation. Network management can be simplified by using port profiles, topology views, and virtualization management.

System Center is used for more advanced levels of management and control, which can significantly reduce deployment and day-to-day maintenance times, while providing in-depth visibility into the network performance and operations of Lenovo switches. Furthermore, when leveraging tools like VMware vCenter Server or vSphere, Switch Center provides additional integration for better optimization.

Features and specifications

The Flex System Fabric CN4093 10Gb Converged Scalable Switch has the following features and specifications:

- **Internal ports**
 - Forty-two internal full-duplex 10 Gigabit ports
 - Two internal full-duplex 1 GbE ports connected to the Chassis Management Module
- **External ports**
 - Two ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or SFP+ direct-attach copper (DAC) cables. SFP+ modules and DAC cables are not included and must be purchased separately.
 - Twelve Omni Ports, each of which can operate as a 10 Gb Ethernet (support for 10GBASE-SR, 10GBASE-LR, or 10 GbE SFP+ DAC cables), or auto-negotiating 4/8 Gb Fibre Channel, depending on the SFP+ transceiver installed in the port. SFP+ modules and DAC cables are not included and must be purchased separately.
Note: Omni Ports do not support 1 Gb SFP Ethernet transceivers.
 - Two ports for 40 Gb Ethernet QSFP+ transceivers or QSFP+ DAC cables. In addition, you can use break-out cables to break out each 40 GbE port into four 10 GbE SFP+ connections. 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 switch module.
- **Scalability and performance**
 - 40 Gb Ethernet ports for more external bandwidth and performance
 - Fixed-speed 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 on Ethernet ports
 - Media access control (MAC) address learning: automatic update, support for up to 128,000 MAC addresses
 - Up to 128 IP interfaces per switch
 - Static and LACP (IEEE 802.3ad) link aggregation, up to 220 Gb of total external bandwidth per switch, up to 64 trunk groups, up to 16 ports per group
 - Support for jumbo frames (up to 9,216 bytes)
 - Broadcast/multicast storm control

- IGMP snooping to limit flooding of IP multicast traffic
- IGMP filtering to control multicast traffic for hosts participating in multicast groups
- Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses, or both
- Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
 - IEEE 802.1D STP for providing L2 redundancy
 - IEEE 802.1s Multiple STP (MSTP) for topology optimization; up to 32 STP instances are supported by a single switch
 - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic such as voice or video
 - Per-VLAN Rapid STP (PVRST) enhancements
 - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on compute nodes
 - Hot Links provides basic link redundancy with fast recovery for network topologies that require Spanning Tree to be turned off
- VLAN support
 - Up to 4095 VLANs supported per switch, with VLAN numbers ranging from 1 to 4095 (4095 is used for management module's connection only.)
 - 802.1Q VLAN tagging support on all ports
 - Private VLANs
- Security
 - VLAN-based, MAC-based, and IP-based access control lists (ACLs)
 - 802.1x port-based authentication
 - 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
 - IPv6 ACL metering
- Quality of Service (QoS)
 - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
 - Traffic shaping and re-marking based on defined policies
 - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic
- IP v4 Layer 3 functions
 - Host management
 - IP forwarding
 - IP filtering with ACLs; up to 896 ACLs supported
 - VRRP for router redundancy
 - Support for up to 128 static routes
 - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP-4); up to 2048 entries in a routing table
 - Support for DHCP Relay
 - Support for IGMP snooping and IGMP relay
 - Support for Protocol Independent Multicast (PIM) in Sparse Mode (PIM-SM) and Dense Mode (PIM-DM)
- IP v6 Layer 3 functions
 - IPv6 host management (except default switch management IP address)
 - IPv6 forwarding
 - Up to 128 static routes
 - Support for OSPF v3 routing protocol
 - IPv6 filtering with ACLs
 - Virtual Station Interface Data Base (VSIDB) support

- Virtualization
 - Virtual NIC (vNIC)
 - 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
 - Virtual link aggregation groups (vLAGs)
 - 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 switch. Data plane traffic is not shared between SPARs on the same switch.
 - 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 switch 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.
 - Multi-hop RDMA over Converged Ethernet (RoCE) with LAG support.
- Fibre Channel and Fibre Channel over Ethernet (FCoE)
 - FC-BB-5 FCoE specification compliant
 - Native FC Forwarder (FCF) switch operations
 - End-to-end FCoE support (initiator to target)
 - FCoE Initialization Protocol (FIP) support
 - FCoE Link Aggregation Group (LAG) support
 - Optimized FCoE to FCoE forwarding
 - Omni Ports support 4/8 Gb FC when FC SFPs+ are installed in these ports
 - Support for F_port, E_Port ISL, NP_port and VF_port FC port types
 - Full Fabric mode for end-to-end FCoE or FCoE gateway; NPV Gateway mode for external FC SAN attachments (support for Brocade and Cisco MDS external SANs)
 - Sixteen buffer credits supported
 - Fabric Device Management Interface (FDMI)
 - NPIV support
 - Fabric Shortest Path First (FSPF)
 - Port security
 - Fibre Channel ping, debugging
 - Supports 2,000 secure FCoE sessions with FIP Snooping by using Class ID ACLs
 - Fabric services in Full Fabric mode:
 - Name Server
 - Registered State Change Notification (RSCN)

- Login services
 - Zoning
- Stacking
 - Hybrid stacking support (from two to six EN4093/EN4093R switches with two CN4093 switches)
 - single IP management
 - FCoE support
 - FCoE LAG on external ports
 - 802.1Qbg support
 - vNIC and UFP support
 - Support for UFP with 802.1Qbg
- Manageability
 - Simple Network Management Protocol (SNMP V1, V2 and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - SSH
 - Secure FTP (sFTP)
 - Service Location Protocol (SLP)
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) for switch clock synchronization
- Monitoring
 - Switch LEDs for external port status and switch module status indication
 - Remote Monitoring (RMON) agent to collect statistics and proactively monitor switch performance
 - Port mirroring for analyzing network traffic passing through the switch
 - Change tracking and remote logging with syslog feature
 - Support for sFLOW agent for monitoring traffic in data networks (separate sFLOW analyzer required elsewhere)
 - POST diagnostics

The following features are not supported with IPv6:

- Default switch management IP address
- SNMP trap host destination IP address
- Bootstrap Protocol (BOOTP) and DHCP
- RADIUS, TACACS+, and LDAP
- QoS metering and re-marking ACLs for out-profile traffic
- VMware Virtual Center (vCenter) for VMready
- Routing Information Protocol (RIP)
- Internet Group Management Protocol (IGMP)
- Border Gateway Protocol (BGP)
- Virtual Router Redundancy Protocol (VRRP)
- sFLOW

Standards supported

The CN4093 switch supports the following standards:

- IEEE 802.1AB Data Center Bridging Capability Exchange Protocol (DCBX)
- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1p Class of Service (CoS) prioritization
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1Qbg Edge Virtual Bridging
- IEEE 802.1Qbb Priority-Based Flow Control (PFC)
- IEEE 802.1Qaz Enhanced Transmission Selection (ETS)
- IEEE 802.1x port-based authentication
- IEEE 802.1w Rapid STP (RSTP)
- 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-KR backplane 10 Gb Ethernet
- 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 802.3ba 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
- FC-PH, Revision 4.3 (ANSI/INCITS 230-1994)
- FC-PH, Amendment 1 (ANSI/INCITS 230-1994/AM1 1996)
- FC-PH, Amendment 2 (ANSI/INCITS 230-1994/AM2-1999)
- FC-PH-2, Revision 7.4 (ANSI/INCITS 297-1997)
- FC-PH-3, Revision 9.4 (ANSI/INCITS 303-1998)
- FC-PI, Revision 13 (ANSI/INCITS 352-2002)
- FC-PI-2, Revision 10 (ANSI/INCITS 404-2006)
- FC-PI-4, Revision 7.0
- FC-FS, Revision 1.9 (ANSI/INCITS 373-2003)
- FC-FS-2, Revision 0.91
- FC-FS-3, Revision 1.11
- FC-LS, Revision 1.2
- FC-SW-2, Revision 5.3 (ANSI/INCITS 355-2001)
- FC-SW-3, Revision 6.6 (ANSI/INCITS 384-2004)
- FC-SW-5, Revision 8.5 (ANSI/INCITS 461-2010)
- FC-GS-3, Revision 7.01 (ANSI/INCITS 348-2001)
- FC-GS-4, Revision 7.91 (ANSI/INCITS 387-2004)
- FC-GS-6, Revision 9.4, (ANSI/INCITS 463-2010)
- FC-BB-5, Revision 2.0 for FCoE
- FCP, Revision 12 (ANSI/INCITS 269-1996)
- FCP-2, Revision 8 (ANSI/INCITS 350-2003)
- FCP-3, Revision 4 (ANSI/INCITS 416-2006)
- FC-MI, Revision 1.92 (INCITS TR-30-2002, except for FL-ports and Class 2)
- FC-MI-2, Revision 2.6 (INCITS TR-39-2005)
- FC-SP, Revision 1.6
- FC-DA, Revision 3.1 (INCITS TR-36-2004)

Chassis and adapters

The switches 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 CN4093 switch 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 CN4093 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 (00D5845) or Upgrade 2 (00D5847) was required for the switch to allow communications on all four ports. With eight-port adapters, both optional Upgrade 1 (90Y3562) and Upgrade 2 (00D5847) were required for the switch 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 switch upgrades for 4-port and 8-port adapters if the total number of port licenses on the switch 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 CN4093 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 CN4093 10Gb Converged Scalable Switch	00D5823	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.

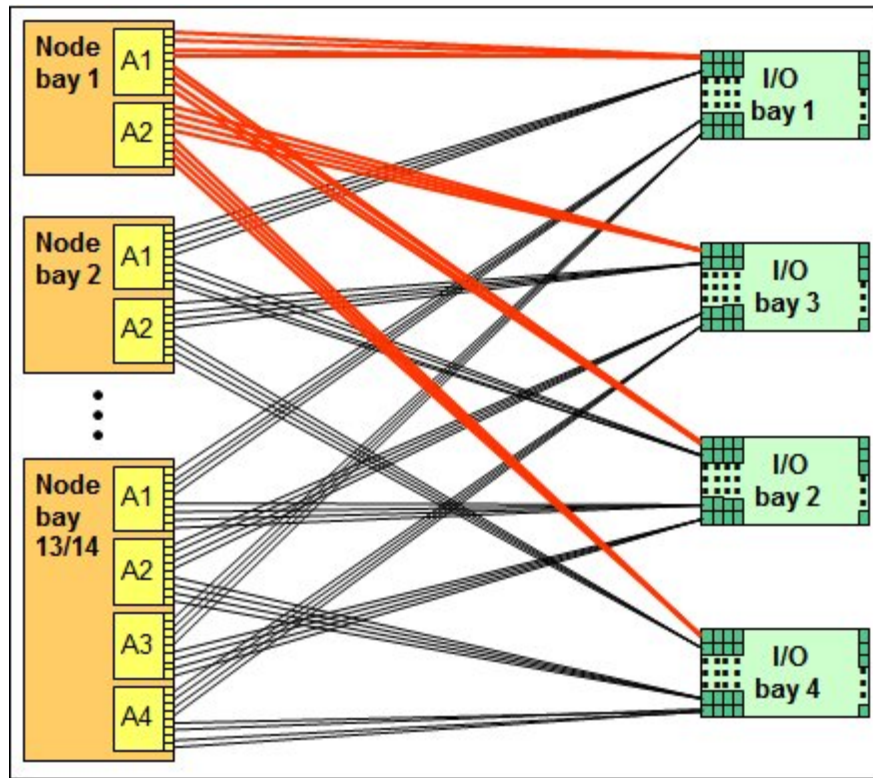


Figure 3. Logical layout of the interconnects between I/O adapters and I/O modules

Table 5 shows the connections between adapters installed in the compute nodes to the switch 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 CN4093.

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 CN4093 10Gb Converged Scalable Switch.




Figure 5. Front panel of the Flex System Fabric CN4093 10Gb Converged Scalable Switch

The following components are located on the front panel:

- LEDs that display the status of the switch module and the network:
 - The OK LED indicates that the switch module has passed the power-on self-test (POST) with no critical faults and is operational.
 - Identify: This blue LED can be used to identify the switch physically by illuminating via the management software.
 - The error LED (switch module error) indicates that the switch module has failed the POST or detected an operational fault.
- One mini-USB RS-232 console port that provides an additional means to configure the switch module. This mini-USB-style connector enables connection of a special serial cable. (The cable is optional and it is not included with the switch. See the Part number information section for details.)
- Two external SFP/SFP+ ports for 1 GbE or 10 GbE connections to external Ethernet devices.
- Twelve external SFP+ Omni Ports for 10 GbE connections to the external Ethernet devices or 4/8 Gb FC connections to the external SAN 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.
- A link OK LED and a Tx/Rx LED for each external port on the switch module.
- A mode LED for each pair of Omni Ports indicating the operating mode. (OFF indicates the port pair is configured for Ethernet operation, and ON indicates the port pair is configured for Fibre Channel operation.)

Network cabling requirements

The network cables that can be used with the CN4093 switch are shown in Table 7.

Table 7. CN4093 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
40GBASE-SR 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 (OM3 multimode fiber) or up to 150 m (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, 00NU537)	10GBASE-SR	Up to 30 m with fiber cables supplied by Lenovo (see Table 3); 850 nm OM3 multimode fiber (50 μ or 62.5 μ) 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
8 Gb Fibre Channel (operating at 8 Gbps or 4 Gbps)			
8Gb SFP+ SW Optical Transceiver (44X1964)	FC-P1-4 (8GFC, 4GFC)	Up to 30 m with fiber optic cables supplied by Lenovo (see Table 3); 850 nm multimode fiber, 50 μ (up to 150 m at 8 Gbps; up to 380 m at 4 Gbps) or 62.5 μ (up to 21 m at 8 Gbps; up to 70 m at 4 Gbps)	LC
Management ports			
External 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 switch carries a 1-year, customer-replaceable unit (CRU) limited warranty. When installed in a chassis, the switch assumes your system's base warranty and any Lenovo warranty service upgrade for the system.

Physical specifications

These are the approximate dimensions and weight of the switch:

- 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 switches conform to the following standards:

- 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 following usage scenarios are described:

- CN4093 FCoE Virtual Fabric in the Full Fabric mode (end-to-end FCoE)
- CN4093 FCoE Virtual Fabric in the NPV Gateway mode (FC Forwarder)
- CN4093 with flexible port mapping (4-port network adapter example)

CN4093 FCoE Virtual Fabric in the Full Fabric mode (end-to-end FCoE)

The CN4093 Virtual Fabric vNIC solution is based on the Flex System Enterprise Chassis with a 10 Gb Converged Enhanced Ethernet (CEE) infrastructure and 10 Gb Virtual Fabric Adapters (VFAs) installed in each compute node. In Virtual Fabric mode, the CN4093 10 Gb switch is vNIC-aware, that is, the configuration of vNICs is done on a switch, then it propagates vNIC parameters to VFA using the DataCenter Bridging eXchange (DCBX) protocol. vNIC bandwidth allocation and metering is performed by both the switch and the VFA. In such a case, a bidirectional virtual channel of an assigned bandwidth is established between them for every defined vNIC. Up to 32 vNICs can be configured on a half-wide compute node.

In the Full Fabric mode, the CN4093 converged switch has 10 GbE external ports to the G8264 top-of-rack switch for external LAN connectivity, and is connected to the integrated Flex System V7000 Storage Node with native FCoE interface via internal 10 GbE links, as shown in Figure 6.




Figure 6. CN4093 Virtual Fabric in the Full Fabric mode (end-to-end FCoE)

The solution components used in the scenario depicted in Figure 6 are listed in Table 8.

Table 8. Components used in an end-to-end FCoE solution with the CN4093 switch (Figure 6)

Diagram reference	Description	Part number	Quantity
1	Flex System end-to-end FCoE solution		
	Flex System x240 Compute Node or other supported compute node	Varies	Varies
	Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per node
	Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	2
	Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)*	00D5845	1 per CN4093
2	Flex System V7000 Storage Node		
3	RackSwitch G8264		

* The Upgrade 1 might not be needed with flexible port mapping if the total number of the internal and external ports used on the CN4093 is less or equal to 22.

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

CN4093 FCoE Virtual Fabric in the NPV Gateway mode

As part of an FCoE solution, the Flex System Fabric CN4093 10Gb Converged Scalable Switch can operate as an integrated FC Forwarder in the NPV Gateway mode, providing the capability to connect to the external Brocade or Cisco MDS storage networks, as shown in Figure 7 and Figure 8.

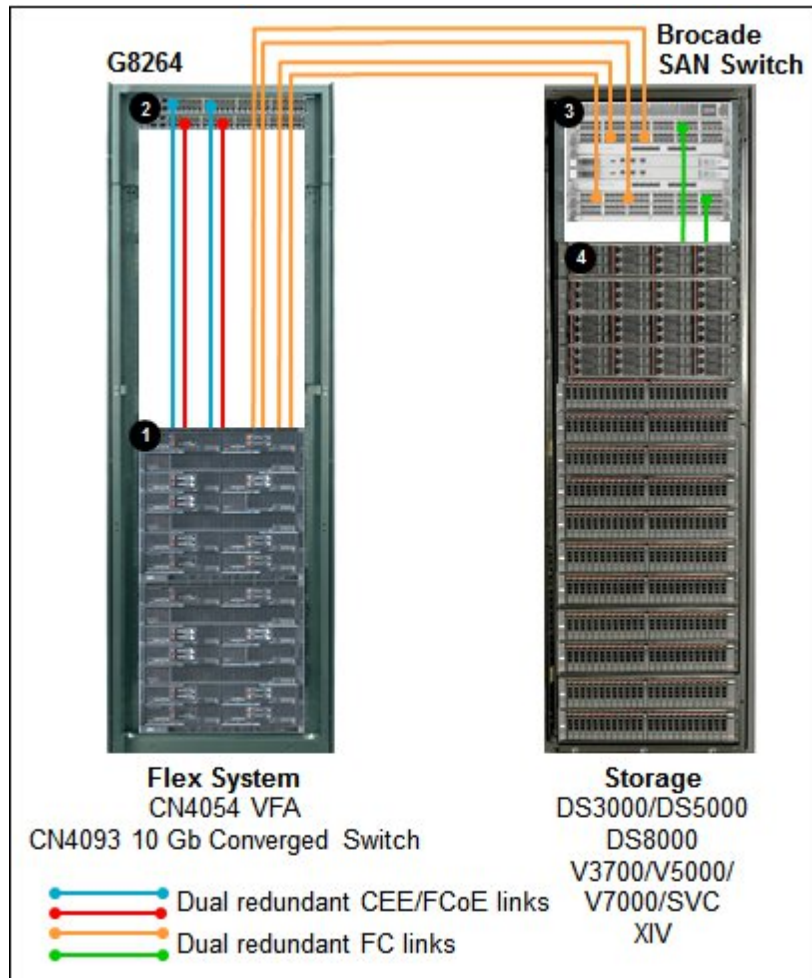


Figure 7. CN4093 as an NPV Gateway connected to the Brocade SAN

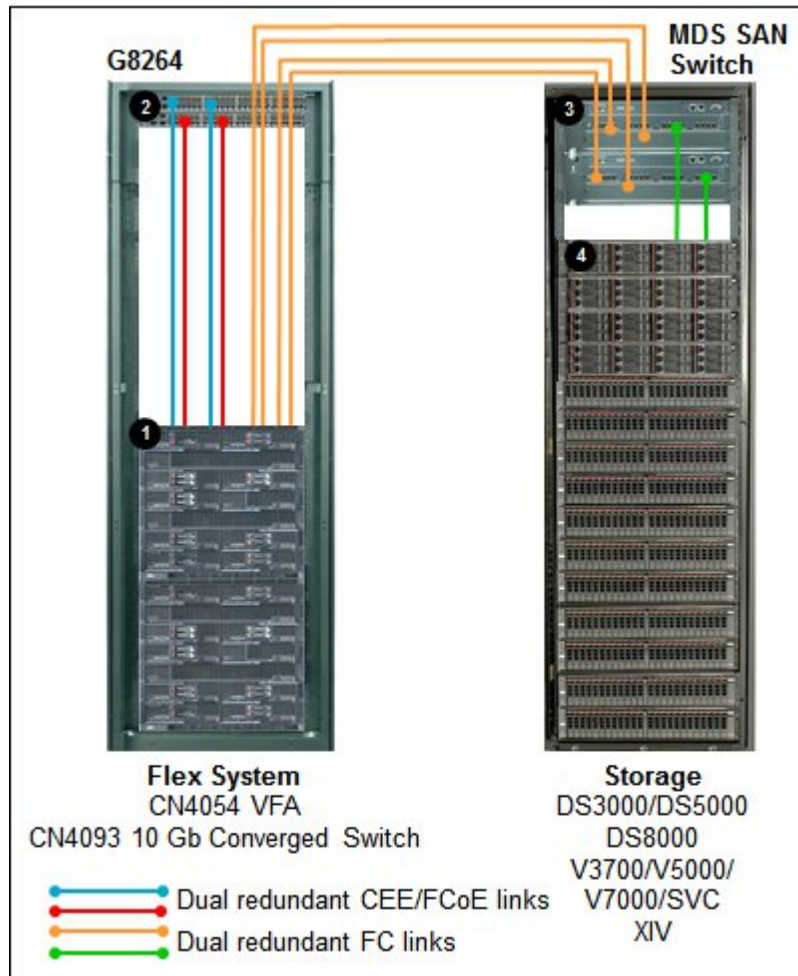


Figure 8. CN4093 as an NPV Gateway connected to the Cisco MDS SAN

The solution components used in the scenarios depicted in Figure 7 and Figure 8 are listed in Table 9 and Table 10 respectively.

Table 9. CN4093 as an NPV Gateway connected to the Brocade SAN (Figure 7)

Diagram reference	Description	Part number	Quantity
1	Flex System FCoE solution		
	Flex System x240 Compute Node or other supported compute node	Varies	Varies
	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 CN4093 10Gb Converged Scalable Switch	00D5823	2
	Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)*	00D5845	1 per CN4093
2	RackSwitch G8264		
3	Brocade 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 CN4093 is less or equal to 22.

Table 10. CN4093 as an NPV Gateway connected to the Cisco MDS SAN (Figure 8)

Diagram reference	Description	Part number	Quantity
1	Flex System FCoE solution		
	Flex System x240 Compute Node or other supported compute node	Varies	Varies
	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 CN4093 10Gb Converged Scalable Switch	00D5823	2
	Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)*	00D5845	1 per CN4093
2	RackSwitch G8264		
3	Cisco MDS 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 CN4093 is less or equal to 22.

Note: You also need SFP+ modules and optical cables or SFP+ DAC cables (not shown in Table 8 and Table 9; see Table 3 for details) for the external 10 Gb Ethernet connectivity, and SW 8 Gb FC SFP+ transceivers and optical cables (also identified in Table 3) for the external Fibre Channel 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>

CN4093 with flexible port mapping (4-port network adapter example)

Prior to Networking OS 7.8, compute nodes with 4-port network adapters required Upgrade 1 or Upgrade 2 for the CN4093 to enable connectivity on all four adapter ports despite of number of compute nodes and external connections used. With the introduction of flexible port mapping in Networking OS 7.8, if the Flex System chassis is not fully populated with the compute nodes that have four network ports, there might be no need to buy Upgrade 1.

Consider the following scenario. You are planning to install nine x240 compute nodes with CN4054 adapters or nine high-density x222 compute nodes that will be connected to two CN4093 switches installed in I/O bays 1 and 2. You are also planning to use two external 10 GbE ports on each CN4093 for the connectivity to the upstream network and two Omni Ports on each switch for FC SAN connectivity. In this scenario, the total number of 10 GbE ports and Omni Ports needed per one CN4093 is 22. The base switch supplies required 22 port licenses; therefore, the solution can be implemented without the need to buy Upgrade 1 or Upgrade 2.

Figure 10 illustrates this scenario.

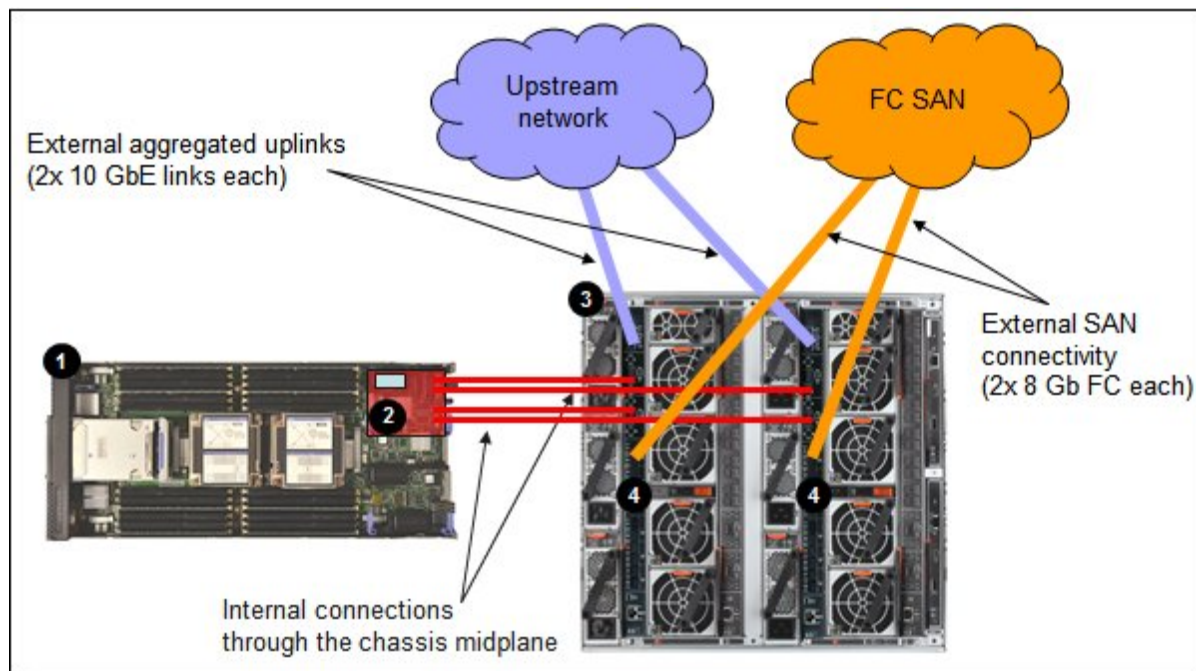


Figure 9. CN4093 with flexible port mapping (4-port network adapter example)

The solution components used in the scenario shown in Figure 10 are listed in Table 11.

Table 11. CN4093 with 4-port network adapters

Diagram reference	Description	Part number	Quantity
1	Flex System x240 Compute Node or other supported compute node	Varies	Up to 9 per chassis
2	Flex System CN4054R 10Gb Virtual Fabric Adapter	00Y3306	1 per compute node
	Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
3	Flex System Enterprise Chassis with additional power supplies and fan modules if needed	8721A1G	1
4	Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	2 per chassis
	SFP+ modules and optical cables or SFP+ DAC cables (See Table 3)	Varies	4 (2 per CN4093)
	8Gb SFP+ SW Optical Transceiver	44X1964	4 (2 per CN4093)

Related publications and links

For more information see the following Flex System Fabric CN4093 10Gb Converged Scalable Switch product publications, available from the Flex System Information Center:

<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

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

These are other useful references:

- US Announcement Letter for the CN4093
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS112-185>
- Flex System Enterprise Chassis Product Guide
<http://lenovopress.com/tips0863>
- *Flex System Products and Technology*, SG24-8255
<http://lenovopress.com/sg248255>
- Flex System Interoperability Guide
<http://lenovopress.com/fsig>
- Support: FixCentral
<http://www.ibm.com/support/fixcentral/>

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 2024. All rights reserved.

This document, TIPS0910, 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/TIPS0910>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS0910>.

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
Omni Ports
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.