IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

IBM Redbooks Product Guide

The IBM Flex System™ Fabric CN4093 10Gb Converged Scalable Switch provides unmatched scalability, 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 IBM Flex System Enterprise Chassis. The switch can help clients migrate to a 10 Gb or 40 Gb converged Ethernet infrastructure and offers virtualization features like Virtual Fabric and VMready®, plus the ability to work with IBM® Distributed Virtual Switch 5000V.

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

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

Did you know

The CN4093 offers up to 12 external programmable IBM 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 CN4093's flexible port licensing allows you to buy only the ports that you need, when you need them. The base switch configuration includes fourteen 10 GbE connections to the node bays, two 10 GbE SFP+ ports, and six Omni Ports with SFP+ connectors. You then have the flexibility of turning on more 10 GbE connections to the server, and more Omni Ports and 40 GbE QSFP+ uplink ports (or 4x 10 GbE SFP+ DAC uplinks on each QSFP+ port) when you need them using IBM Features on Demand licensing capabilities that provide “pay as you grow” scalability without the need for additional hardware.

Part number information

The CN4093 switch is initially licensed for fourteen 10 GbE internal ports, two external 10 GbE SFP+ ports, and six external Omni Ports enabled. Further ports can be enabled, including fourteen additional internal ports and two external 40 GbE QSFP+ uplink ports with Upgrade 1, and fourteen additional internal ports and six additional external Omni Ports with the 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 (x-config / e-config)
Switch module		
IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	A3HH / ESW2
Features on Demand upgrades		
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845	A3HL / ESU1
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 2)	00D5847	A3HM / ESU2

The part number for the switch includes the following items:

- One IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch
- Important Notices Flyer
- Warranty Flyer
- Technical Update Flyer
- Documentation CD-ROM

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, IBM 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:

- Feature on Demand Activation Flyer
- Upgrade activation key

The base switch and upgrades are as follows:

- 00D5823 is the part number for the physical device, which comes with 14 internal 10 GbE ports enabled (one to each node bay), two external 10 GbE SFP+ ports enabled to connect to a top-of-rack switch or other devices, and six Omni Ports enabled to connect to either Ethernet or Fibre Channel networking infrastructure, depending on the SFP+ cable or transceiver used.
- 00D5845 (Upgrade 1) can be applied on the base switch when you need more uplink bandwidth with two 40 GbE QSFP+ ports that can be converted into 4x 10 GbE SFP+ DAC links with the optional break-out cables. This upgrade also enables 14 additional internal ports, for a total of 28 ports, to provide more bandwidth to the compute nodes leveraging 4-port expansion cards.
- 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 Gb ports, for a total of 28 internal ports, to provide more bandwidth to the compute nodes leveraging four-port expansion cards.
- Both 00D5845 (Upgrade 1) and 00D5847 (Upgrade 2) can be applied on the switch at the same time to allow you to leverage six ports on an eight-port expansion card, and to utilize all external ports on the switch.

The following table lists supported port combinations on the switch and required upgrades.

Table 2. Supported port combinations

	Quantity required		
Supported port combinations	**Base switch, 00D5823**	**Upgrade 1, 00D5845**	**Upgrade 2, 00D5847**
- 14x internal 10 GbE ports - 2x external 10 GbE SFP+ ports - 6x external SFP+ Omni Ports	1	0	0
- 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
- 28x internal 10 GbE ports - 2x external 10 GbE SFP+ ports - 12x external SFP+ Omni Ports	1	0	1
- 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.

Supported transceivers and cables

Table 3 lists the supported cables and transceivers.

Table 3. Supported transceivers and direct-attach cables

Description	Part number	Feature code (x-config / e-config)
Serial console cables		
IBM Flex System Management Serial Access Cable Kit	90Y9338	A2RR / None
SFP transceivers - 1 GbE (supported on two dedicated SFP+ ports)		
IBM SFP RJ-45 Transceiver (does not support 10/100 Mbps)	81Y1618	3268 / EB29
IBM SFP SX Transceiver	81Y1622	3269 / EB2A
IBM SFP LX Transceiver	90Y9424	A1PN / ECB8
SFP+ transceivers - 10 GbE (supported on SFP+ ports and Omni Ports)		
IBM SFP+ SR Transceiver	46C3447	5053 / EB28
IBM SFP+ LR Transceiver	90Y9412	A1PM / ECB9
10GBase-SR SFP+ (MMFiber) transceiver	44W4408	4942 / 3382
SFP+ direct-attach cables - 10 GbE (supported on SFP+ ports and Omni Ports)		
1m IBM Passive DAC SFP+	90Y9427	A1PH / ECB4
3m IBM Passive DAC SFP+	90Y9430	A1PJ / ECB5
5m IBM Passive DAC SFP+	90Y9433	A1PK / ECB6
QSFP+ transceiver and cables - 40 GbE (supported on QSFP+ ports)		
IBM QSFP+ 40GBASE-SR Transceiver (Requires either cable 90Y3519 or cable 90Y3521)	49Y7884	A1DR / EB27
10m IBM MTP Fiber Optical Cable (requires transceiver 49Y7884)	90Y3519	A1MM / EB2J
30m IBM MTP Fiber Optical Cable (requires transceiver 49Y7884)	90Y3521	A1MN / EC2K
QSFP+ breakout cables - 40 GbE to 4x10 GbE (supported on QSFP+ ports)		
1m 40Gb QSFP+ to 4 x 10Gb SFP+ Cable	49Y7886	A1DL / EB24
3m 40Gb QSFP+ to 4 x 10Gb SFP+ Cable	49Y7887	A1DM / EB25
5m 40Gb QSFP+ to 4 x 10Gb SFP+ Cable	49Y7888	A1DN / EB26
QSFP+ direct-attach cables - 40 GbE (supported on QSFP+ ports)		
1m QSFP+ to QSFP+ DAC	49Y7890	A1DP / EB2B
3m QSFP+ to QSFP+ DAC	49Y7891	A1DQ / EB2H
SFP+ transceivers - 8 Gb FC (supported on Omni Ports)		
IBM 8Gb SFP+ SW Optical Transceiver	44X1964	5075 / 3286

With the flexibility of the CN4093 switch, clients can take advantage of the technologies that they require for multiple environments:

- For 1 GbE links (supported on external SFP+ ports 1 and 2 only), you can use 1 GbE SFP transceivers plus RJ-45 cables or LC-to-LC fiber cables, depending on the transceiver.
- For 10 GbE (supported on external SFP+ ports and Omni Ports), you can use direct-attached cables (DAC, also known as Twinax), which come in lengths between 1 m and 5 m. These DACs are a cost-effective and low-power alternative to transceivers, and are ideal for all 10 Gb Ethernet connectivity within the rack, or even connecting to an adjacent rack. For longer distances, there is a choice of SFP+ transceivers (SR or LR) plus LC-to-LC fiber optic cables.
- For 40 GbE links (supported on QSFP+ ports), you can use QSFP+ to QSFP+ cables up to 3 m, or QSFP+ transceivers and MTP cables for longer distances. You also have the ability to break out the 40 GbE ports into four 10 GbE connections using optional break-out cables.
- For 8 Gb FC links (supported on Omni Ports only), you can use SFP+ SW transceivers plus LC-to-LC fiber optics cables.

Benefits

The IBM 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 uplinks.
- 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 flexibility

  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; IBM Omni Ports give clients the flexibility to choose between 10 Gb Ethernet uplink connections to the top-of-rack switch and 4/8 Gb Fibre Channel for flexible, scalable, and convenient access to FC storage.

- Increased network 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 server 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 uplink bandwidth by being the integrated switch to support 40 Gb uplinks.

- Pay as you grow flexibility

  The CN4093's flexible port licensing allows you to buy only the ports that you need, when you need them. The base switch configuration includes fourteen 10 GbE connections to the servers, two 10 GbE uplinks, and six Omni Ports for either 10 GbE or 4/8 Gb FC connections. You then have the flexibility of turning on more 10 GbE connections to the server and more Omni Ports or 40 GbE uplinks when you need them using Features on Demand licensing capabilities, which provide pay as you grow scalability without the need for additional hardware.

- 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. IBM is helping to address these requirements by removing multiple physical I/O ports. IBM Virtual Fabric provides a way for companies to carve up 10 Gb ports into virtual NICs (vNICs) to meet those requirements with Intel processor-based servers.

  To help deliver maximum performance per vNIC and to provide higher availability and security with isolation between vNICs, the switch leverages capabilities of its IBM Networking Operating System. For large-scale virtualization, the IBM 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.

- VM-aware networking

  Delivering advanced virtualization awareness helps simplify management and automates VM mobility by making the network VM aware with IBM VMready, which works with all the major hypervisors. For companies using VMware, IBM System Networking's Distributed Virtual Switch 5000V (sold separately) enables network administrators to simplify management by having a consistent virtual and physical networking environment. 5000V virtual and physical servers 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 server to server.

  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 server 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.

- Integrate network management

  A key challenge is the management of a discrete network environment. The CN4093 10Gb Converged Scalable Switch is tightly integrated and managed through the IBM Flex System Manager. The switch also 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.

  For more advanced levels of management and control, IBM System Networking Element Manager (SNEM) can significantly reduce deployment and day-to-day maintenance times, while providing in-depth visibility into the network performance and operations of IBM switches. Furthermore, when leveraging tools like VMware vCenter Server (formerly VMware VirtualCenter) or vSphere, SNEM provides additional integration for better optimization.

Features and specifications

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

- Internal ports
  - Forty-two internal full-duplex 10 Gigabit ports (Fourteen ports are enabled by default. Optional FoD licenses are required to activate the remaining 28 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+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or SFP+ copper direct-attach cables (DACs)). These two ports are enabled by default. SFP+ modules and DACs are not included and must be purchased separately.
  - Twelve IBM Omni Ports, each of which can operate as a 10 Gb Ethernet (support for 10GBASE-SR, 10GBASE-LR, or 10 GbE SFP+ DACs), or auto-negotiating 4/8 Gb Fibre Channel, depending on the SFP+ transceiver installed in the port. The first six ports are enabled by default. An optional FoD license is required to activate the remaining six ports. SFP+ modules and DACs are not included and must be purchased separately.

    Note: Omni Ports do not support 1 Gb Ethernet operations.
  - Two ports for 40 Gb Ethernet QSFP+ transceivers or QSFP+ DACs (Ports are disabled by default. An optional FoD license is required to activate them). In addition, you can use break-out cables to break out each 40 GbE port into four 10 GbE SFP+ connections. QSFP+ modules and DACs 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 extreme uplink bandwidth and performance
  - Fixed-speed external 10 Gb Ethernet ports to leverage 10 Gb core 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 uplink 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 1024 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
- 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

- Virtualization
  - Virtual NICs (vNICs)
    - Ethernet, iSCSI, or FCoE traffic is supported on vNICs
  - 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

- 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.

- Fibre Channel over Ethernet (FCoE)
  - FC-BB5 FCoE specification compliant
  - Native FC Forwarder switch operations
  - End-to-end FCoE support (initiator to target)
  - FCoE Initialization Protocol (FIP) support

- Fibre Channel
  - Omni Ports support 4/8 Gb FC when FC SFPs+ are installed in these ports
  - Full Fabric mode for end-to-end FCoE or NPV Gateway mode for external FC SAN attachments (support for IBM B-type, Brocade, and Cisco MDS external SANs)
  - Fabric services in Full Fabric mode:
    - Name Server
    - Registered State Change Notification (RSCN)
    - Login services
    - Zoning

- 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 CN4058 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 802.2 Logical Link Control
- 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 802.3ba 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-BB-5 FCoE

Supported chassis and adapter cards

The switches are installed in switch bays in the rear of the IBM Flex System Enterprise Chassis, as shown in the following figure. Switches are normally installed in pairs because I/O adapter cards installed in the compute nodes route to two switch bays for redundancy and performance.

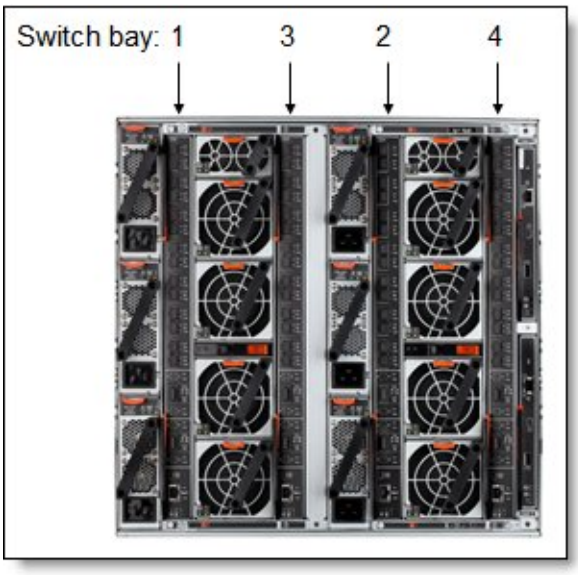

Figure 2. Location of the switch bays in the IBM Flex System Enterprise Chassis

The 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 servers, such as the x240 with two adapters, and full-wide servers, such as the p460 with four adapters.

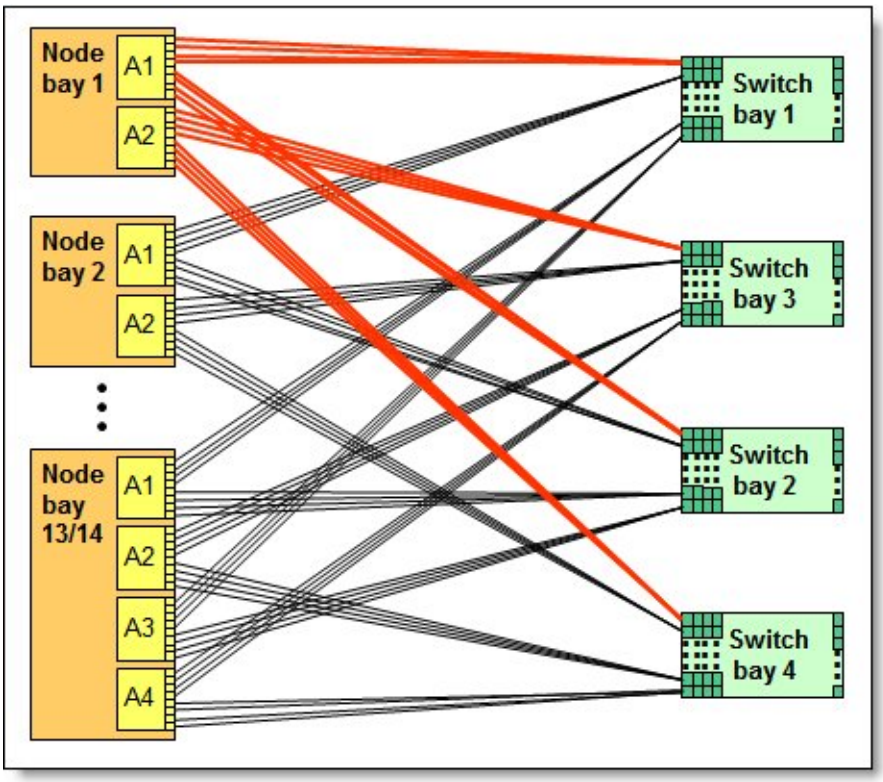

Figure 3. Logical layout of the interconnects between I/O adapters and I/O modules

The CN4093 switch can be installed in bays 1, 2, 3, and 4 of the Enterprise chassis. A supported adapter card must be installed in the corresponding slot of the compute node (slot A1 when switches are installed in bays 1 and 2 or slot A2 when switches are in bays 3 and 4). With four-port adapters, an optional Upgrade 1 (00D5845) or Upgrade 2 (00D5847) is required for the switch to allow communications on all four ports. With eight-port adapters, both optional Upgrade 1 (90Y3562) and Upgrade 2 (88Y6037) are required for the switch to allow communications on six adapter ports. Two remaining ports are reserved for future use.

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 in slot A1 is not required. However, when needed, the periscope connector can be replaced with the adapter card. In this case, integrated NIC will be disabled.

The following table shows the connections between adapters installed in the compute nodes to the switch bays in the chassis.

Table 4. Adapter to I/O bay correspondence

I/O adapter slot in the server	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#	Yes			
	Port 8#		Yes		
Slot 2	Port 1			Yes	
	Port 2				Yes
	Port 3*			Yes	
	Port 4*				Yes
	Port 5**			Yes	
	Port 6**				Yes
	Port 7#			Yes	
	Port 8#				Yes
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#	Yes			
	Port 8#		Yes		
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#			Yes	
	Port 8#				Yes

* Ports 3 and 4 require Upgrade 1 or Upgrade 2 for the CN4093 switch.
** Ports 5 and 6 require both Upgrade 1 and Upgrade 2 for the CN4093 switch.
** Ports 7 and 8 are reserved for future use.

The following table lists the I/O adapters supported by the CN4093 switch.

Table 5. Supported network adapters

Description	Part number	Feature code (x-config / e-config)	CN4093
10 Gb Ethernet			
Embedded 10Gb Virtual Fabric Adapter (2-port)	None	None / None	Yes*
IBM Flex System Embedded 10Gb Virtual Fabric Upgrade (Feature on Demand to provide FCoE and iSCSI support)	90Y9310	A2TD / None	Yes
IBM Flex System CN4054 10Gb Virtual Fabric Adapter (4-port)	90Y3554	A1R1 / None	Yes
IBM Flex System CN4054 Virtual Fabric Adapter (SW Upgrade) (Feature on Demand to provide FCoE and iSCSI support)	90Y3558	A1R0 / None	Yes
IBM Flex System CN4058 8-port 10Gb Converged Adapter	None	None / EC24	Yes
IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	None	None / 1762	Yes
IBM Flex System EN4132 2-port 10Gb Ethernet Adapter	90Y3466	A1QY / None	No
IBM Flex System EN4132 2-port 10Gb RoCE Adapter	None	None / EC26	No
1 Gb Ethernet			
Embedded 1 Gb Ethernet controller (2-port)	None	None / None	Yes
IBM Flex System EN2024 4-port 1Gb Ethernet Adapter	49Y7900	A10Y / 1763	Yes

* The Embedded 10Gb Virtual Fabric Adapter is included in models of the x240 with model numbers of the form x2x.

The adapters are installed in slots in each compute node. The following figure shows the locations of the slots in the x240 Compute Node. The positions of the adapters in the other supported servers are similar.

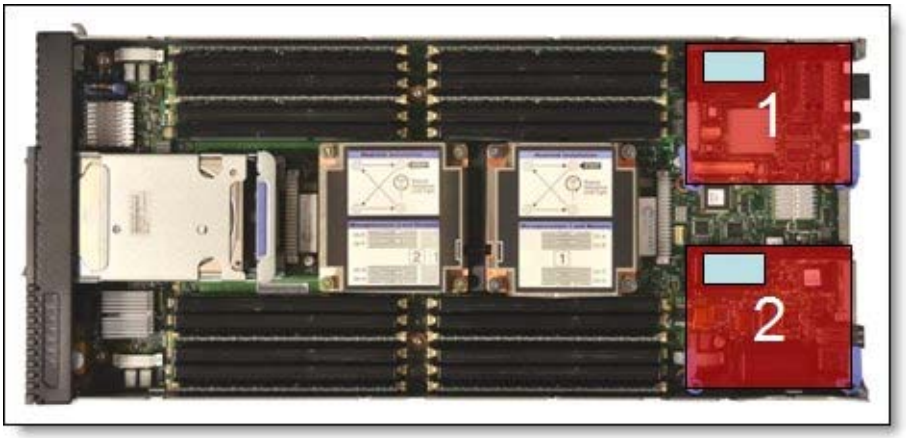

Figure 4. Location of the I/O adapter slots in the IBM Flex System x240 Compute Node

Connectors and LEDs

Figure 5 shows the front panel of the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch.

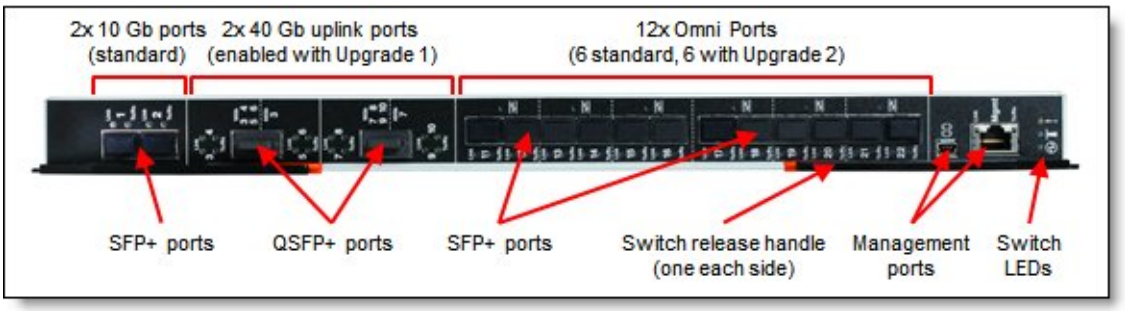

Figure 5. Front panel of the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

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+ ports for 1 Gb or 10 Gb connections to external Ethernet devices.
- Twelve external SFP+ Omni Ports for 10 Gb 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+ modules or cables for a single 40 Gb uplink per port or for splitting of a single port into 4x 10 Gb connections to external Ethernet devices.
- 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 the following table.

Table 6. CN4093 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
IBM QSFP+ 40GBASE-SR Transceiver (49Y7884)	40GBASE-SR4	10 m or 30 m IBM MTP fiber optics cables (see Table 3)	MTP
Direct attach cable	40GBASE-CR4	1 m or 3 m QSFP+ to QSFP+ DACs (see Table 3)	QSFP+
10 Gb Ethernet			
IBM SFP+ SR Transceiver (46C3447)	10GBASE-SR	850 nm multimode fiber cable (50 µ or 62.5 µ) up to 300 m	LC
IBM SFP+ LR Transceiver (90Y9412)	10GBASE-LR	1310 nm single-mode fiber cable up to 10 km	LC
10GBase-SR SFP+ (MMFiber) transceiver (44W4408)	10GBASE-SR	850 nm multimode fiber cable (50 µ or 62.5 µ) up to 300 m	LC
Direct attach cable	10GSFP+Cu	Up to 5 m SFP+ copper DACs (see Table 3)	SFP+
1 Gb Ethernet			
IBM SFP RJ-45 Transceiver (81Y1618)	1000BASE-T	UTP Category 5, 5E, and 6 up to 100 meters	RJ-45
IBM SFP SX Transceiver (81Y1622)	1000BASE-SX	850 nm multimode fiber cable (50 µ or 62.5 µ) up to 550 m	LC
IBM SFP LX Transceiver (90Y9424)	1000BASE-LX	1310 nm single-mode fiber cable up to 10 km	LC
8 Gb Fibre Channel			
IBM 8Gb SFP+ SW Optical Transceiver (44X1964)	FC-PI-4 (8GFC)	850 nm multimode fiber, 50 µ (up to 150 meters) or 62.5 µ (up to 21 meters)	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

There is a 1-year, customer-replaceable unit (CRU) limited warranty. When installed in a chassis, the switch assumes your system's base warranty and any IBM ServicePac® upgrade.

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)

Regulatory compliance

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 Communique 2004/9; Communique 2004/22
- Saudi Arabia EMC.CVG, 28 October 2002

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

The CN4093 Virtual Fabric vNIC solution is based on the IBM 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 uplinks to the G8264 top-of-rack switch for external LAN connectivity, and is connected to the integrated IBM Flex System V7000 Storage Node with native FCoE interface via internal 10 GbE links, as shown in the following figure.

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 7.

Table 7. Components used in an end-to-end FCoE solution with the CN4093 switch (Figure 6)

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

CN4093 FCoE Virtual Fabric in the NPV Gateway mode

As part of an IBM FCoE solution, the IBM 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 IBM B-type, Brocade, or Cisco MDS storage networks, as shown in Figure 7 and Figure 8.

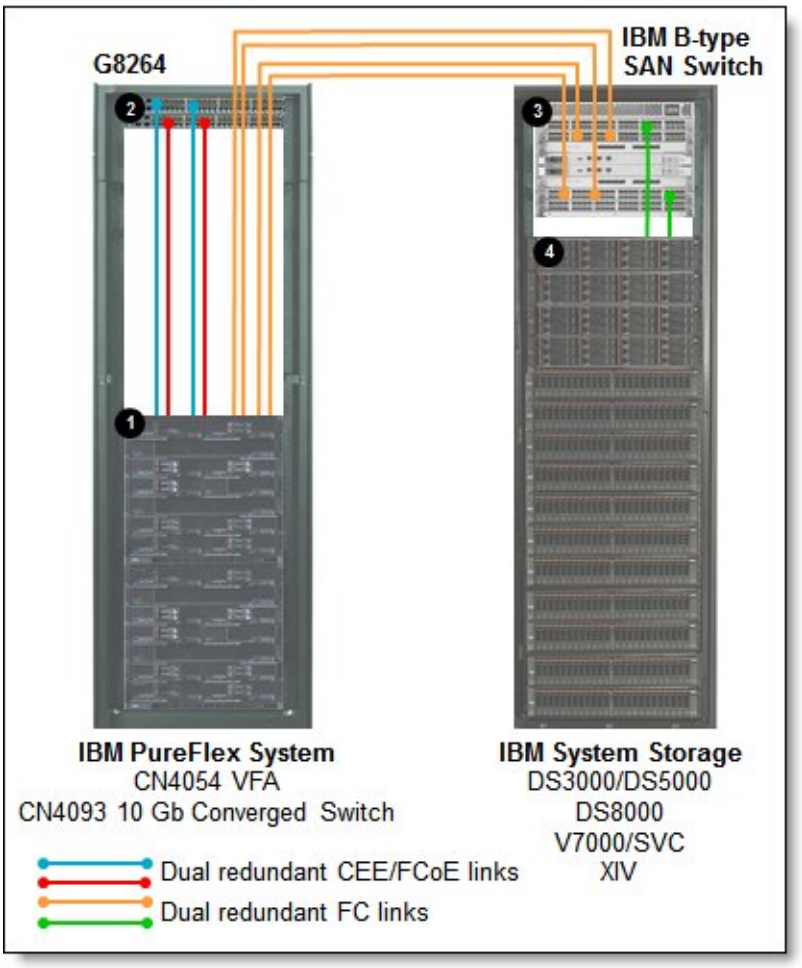

Figure 7. CN4093 as an NPV Gateway connected to the IBM B-type SAN

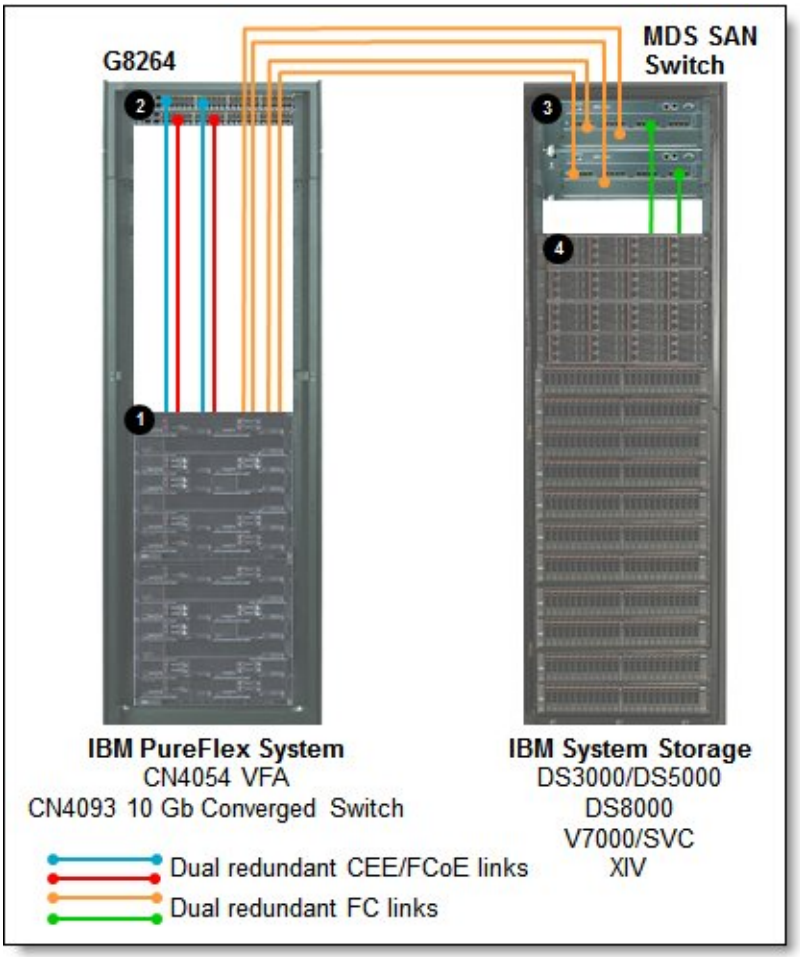

Figure 8. CN4093 as an NPV Gateway connected to the Cisco MDS SAN

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

Table 8. CN4093 as an NPV Gateway connected to the IBM B-type SAN (Figure 7)

Diagram reference	Description	Part number	Quantity
1	**IBM Flex System FCoE solution**		
	IBM Flex System x240 Compute Node or other supported server	Varies	Varies
	IBM Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per server
	IBM Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	2
	IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845	1 per CN4093
2	**IBM RackSwitch G8264**		
3	**IBM B-type or Brocade SAN fabric**		
4	**IBM System Storage FC disk controllers**		
	IBM System Storage DS3000 / DS5000		
	IBM System Storage DS8000		
	IBM Storwize V7000 / SAN Volume Controller		
	IBM XIV		

Table 9. CN4093 as an NPV Gateway connected to the Cisco MDS SAN (Figure 8)

Diagram reference	Description	Part number	Quantity
1	**IBM Flex System FCoE solution**		
	IBM Flex System x240 Compute Node or other supported server	Varies	Varies
	IBM Flex System CN4054 10Gb Virtual Fabric Adapter	90Y3554	1 per server
	IBM Flex System CN4054 Virtual Fabric Adapter Upgrade	90Y3558	1 per VFA
	IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	2
	IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845	1 per CN4093
2	**IBM RackSwitch G8264**		
3	**Cisco MDS SAN fabric**		
4	**IBM System Storage FC disk controllers**		
	IBM System Storage DS3000 / DS5000		
	IBM System Storage DS8000		
	IBM Storwize V7000 / SAN Volume Controller		
	IBM XIV		

Note: You also need SFP+ modules and optical cables or SFP+ DACs (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.

IBM provides extensive FCoE testing to deliver network interoperability. For a full listing of IBM supported FCoE and iSCSI configurations, see the System Storage Interoperation Center (SSIC) website at: http://ibm.com/systems/support/storage/ssic

Related publications

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

- IBM US Announcement Letter for the CN4093
  http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS112-185
- IBM Flex System Enterprise Chassis Product Guide
  http://www.redbooks.ibm.com/abstracts/tips0865.html
- IBM Redbooks® publication *IBM Flex System Products and Technology*, SG24-7984
  http://www.redbooks.ibm.com/abstracts/sg247984.html
- IBM Flex System Interoperability Guide
  http://www.redbooks.ibm.com/abstracts/redpfsig.html?Open
- IBM Redbooks Product Guides for IBM Flex System servers and options
  http://www.redbooks.ibm.com/portals/puresystems
- Configuration and Option Guide
  http://www.ibm.com/systems/xbc/cog/
- IBM Support: FixCentral
  http://www.ibm.com/support/fixcentral/

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM 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:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION 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 states 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM 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 IBM product and use of those Web sites is at your own risk.IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results 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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document was created or updated on November 9, 2012.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips0910.html .

Trademarks