

IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches

IBM Redbooks Product Guide

The IBM Flex System™ Fabric EN4093 and EN4093R 10Gb Scalable Switches provide unmatched scalability and performance, while also delivering innovations to help address a number of networking concerns today and providing capabilities that will help you prepare for the future. These switches are capable of supporting up to sixty-four 10 Gb Ethernet connections while offering Layer 2/3 switching. They are designed to install within the I/O module bays of the IBM Flex System Enterprise Chassis. These switches can help clients migrate to a 10 Gb or 40 Gb Ethernet infrastructure and offer virtualization features like Virtual Fabric and VMready®, plus the ability to work with IBM® Distributed Virtual Switch 5000V.



Figure 1. IBM Flex System Fabric EN4093/EN4093R 10Gb Scalable Switch

Did you know

Flexible port licensing for the EN4093/EN4093R 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 and ten 10 GbE uplinks. You then have the flexibility of turning on more 10 GbE connections to the server and more 10 GbE or 40 GbE uplinks when you need them using IBM Feature on Demand licensing capabilities that provide “pay as you grow” scalability.

IBM Flex System, a new category of computing and the next generation of Smarter Computing, offers intelligent workload deployment and management for maximum business agility. This chassis delivers high-speed performance complete with integrated servers, storage, and networking for multi-chassis management in data center compute environments. Furthermore, its flexible design can meet the needs of varying workloads with independently scalable IT resource pools for higher utilization and lower cost per workload. Increased security and resiliency protect vital information and promote maximum uptime, and the integrated, easy-to-use management system reduces setup time and complexity, providing a quicker path to ROI.

Part number information

The EN4093 and EN4093R switches are initially licensed for fourteen 10 Gb internal ports enabled and ten 10 Gb external uplink ports enabled. Further ports can be enabled, including 14 additional internal ports and two 40 Gb external uplink ports with Upgrade 1, and 14 additional internal ports and four additional SFP+ 10 Gb external ports with Upgrade 2 license options. Upgrade 1 must be applied before Upgrade 2 can be applied. 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 modules		
IBM Flex System Fabric EN4093 10Gb Scalable Switch	49Y4270	A0TB / 3593
IBM Flex System Fabric EN4093R 10Gb Scalable Switch	95Y3309	A3J6 / ESW7
Features on Demand upgrades		
IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	49Y4798	A1EL / 3596
IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 2)	88Y6037	A1EM / 3597

The part number for the switch includes the following items:

- One IBM Flex System Fabric EN4093 or EN4093R 10Gb Scalable Switch
- Important Notices Flyer
- Warranty Flyer
- Documentation CD-ROM

Note: Neither SFP nor SFP+ (small form-factor pluggable plus) 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, 49Y4798 and 88Y6037, include the following items:

- Feature on Demand Activation Flyer
- Upgrade activation key

The base switch and upgrades are as follows:

- 49Y4270 and 95Y3309 are the part numbers for the physical device and come with fourteen internal 10 Gb ports enabled (one to each compute node) and ten external 10 Gb ports enabled to connect to a top-of-rack switch or other 10 Gb devices or to form a switch stacking group. All external 10 Gb ports are SFP+ based connections.
- 49Y4798 (Upgrade 1) can be applied on the base switch when you take full advantage of four-port adapter cards installed in each compute node. This upgrade enables 14 additional internal ports, for a total of 28 ports. The upgrade also enables two 40 Gb uplinks with QSFP+ connectors that can also be used for stacking. These QSFP+ ports can also be converted to four 10 Gb SFP+ DAC connections using the appropriate fan-out cable. This upgrade requires the base switch.
- 88Y6037 (Upgrade 2) can be applied on top of Upgrade 1 when you want more uplink bandwidth on the switch or if you want additional internal bandwidth to the compute nodes with the six-port capable adapter cards. The upgrade will enable the remaining four external 10 Gb uplinks with SFP+ connectors, plus fourteen additional internal 10 Gb ports, for a total of 42 ports (three to each compute node).

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

Table 2. Supported port combinations

Supported port combinations	Quantity required		
	Base switch, 49Y4270 or 95Y3309	Upgrade 1, 49Y4798	Upgrade 2, 88Y6037
<ul style="list-style-type: none"> ● 14x internal 10 GbE ● 10x external 10 GbE 	1	0	0
<ul style="list-style-type: none"> ● 28x internal 10 GbE ● 10x external 10 GbE ● 2x external 40 GbE 	1	1	0
<ul style="list-style-type: none"> ● 42x internal 10 GbE† ● 14x external 10 GbE ● 2x external 40 GbE 	1	1	1

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

Supported cables and transceivers

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		
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		
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		
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		
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		
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		
1m QSFP+ to QSFP+ DAC	49Y7890	A1DP / EB2B
3m QSFP+ to QSFP+ DAC	49Y7891	A1DQ / EB2H

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

- For 1 GbE links, you can use SFP transceivers plus RJ-45 cables or LC-to-LC fiber cables depending on the transceiver.
- For 10 GbE, 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 Gb links, you can use QSFP to QSFP cables up to 3 m, or QSFP transceivers and MTP cables for longer distances. Clients also have the ability to break out the 40 Gb ports into four 10 GbE connections using break-out cables.

Benefits

The IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches are considered particularly suited for these clients:

- Clients who want to use 10 GbE communications between servers in the chassis but still require upstream 1 GbE connections to their existing infrastructure.
- 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.
- Clients who want to implement a converged infrastructure with NAS, iSCSI or FCoE. For FCoE implementations, the EN4093R acts as a transit switch forwarding FCoE traffic upstream to other devices like the Brocade VDX or Cisco Nexus 5548/5596, where the FC traffic is broken out.

The switches offer the following key features and benefits:

- **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 EN4093 and EN4093R are the embedded 10 GbE switches for a server chassis to support sub-microsecond latency and up to 1.28 Tbps, while also delivering full line rate performance, making them ideal for managing dynamic workloads across the network. Furthermore, these switches provide a rich Layer 2 and Layer 3 feature set that is ideal for many of today's data centers, and they offer industry-leading uplink bandwidth by being the first integrated switches to support 40 Gb uplinks.
- **Pay as you grow flexibility**

The EN4093/EN4093R 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 and ten 10 GbE uplinks. You then have the flexibility of turning on more 10 GbE connections to the server and more 10 GbE or 40 GbE uplinks when you need them using Features on Demand licensing capabilities that 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 using a pair of CN4054 10Gb Virtual Fabric Adapters in each compute node and four EN4093 or EN4093R 10Gb Scalable Switches in the chassis.

The EN4093 and EN4093R offer the benefits of IBM's 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.

- 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.

- Simplified network infrastructure

The EN4093 and EN4093R 10Gb Scalable Switches simplify deployment and growth because of their 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 EN4093 and EN4093R support the newest protocols, including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) that can be leveraged in either an iSCSI, Fibre Channel over Ethernet (FCoE), or NAS converged environment.

EN4093R's stacking capabilities simplify management for clients by stacking up to eight 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.

- Integrated network management

A key challenge is the management of a discrete network environment. EN4093 and EN4093R 10Gb Scalable Switches are tightly integrated and managed through the IBM Flex System Manager. The switches also support 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 offers IBM System Networking Element Manager (SNEM) that can significantly reduce deployment and day-to-day maintenance times, while providing in-depth visibility into the network performance and operations of IBM switches. Plus, 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 EN4093 and EN4093R 10Gb Scalable Switches have the following features and specifications:

- Internal ports
 - 42 internal full-duplex 10 Gigabit ports. (14 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
 - 14 ports for 1 Gb or 10 Gb Ethernet SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, or 10GBASE-LR) or SFP+ copper direct-attach cables (DAC). Ten ports are enabled by default. An optional FoD license is required to activate the remaining four ports. SFP+ modules and DACs are not included and must be purchased separately.
 - 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). 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
 - 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 like 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).

- IPv6 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
 - Unified fabric ports (UFPs)
 - Ethernet or FCoE traffic is supported on UFPs
 - 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)
- 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
 - FCoE transit switch operations
 - FCoE Initialization Protocol (FIP) support for automatic ACL configuration
- Stacking
 - Up to eight switches in a stack
 - FCoE support (EN4093R only)
 - vNIC support
 - 802.1Qbg support

- 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) and Precision Time Protocol (PTP) 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 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

The following table summarizes key features supported by the EN4093 and EN4093R.

Table 4. EN4093 and EN4093R supported features

Feature	EN4093	EN4093R
Layer 2 switching	Yes	Yes
Layer 3 switching	Yes	Yes
Switch stacking	Yes	Yes
Virtual NIC (standalone)	Yes	Yes
Virtual NIC (stacking)	Yes	Yes
Unified fabric port (standalone)	Yes	Yes
Unified fabric port (stacking)	No	No
Edge virtual bridging (standalone)	Yes	Yes
Edge virtual bridging (stacking)	Yes	Yes
FCoE (standalone)	Yes	Yes
FCoE (stacking)	No	Yes

Standards supported

The switches support 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

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.

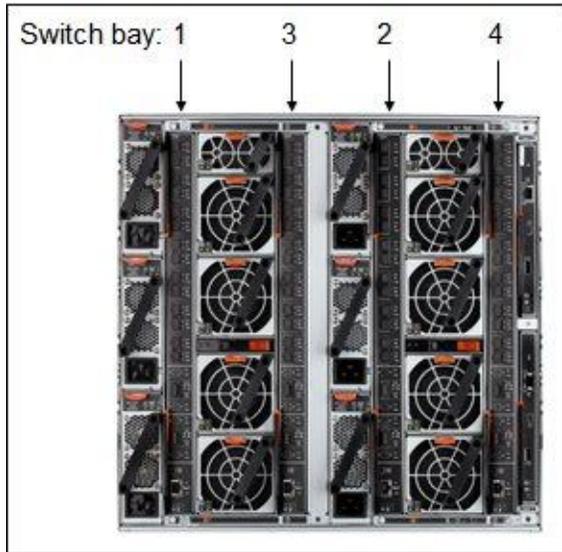


Figure 2. Location of the switch bays in the IBM Flex System Enterprise Chassis

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.

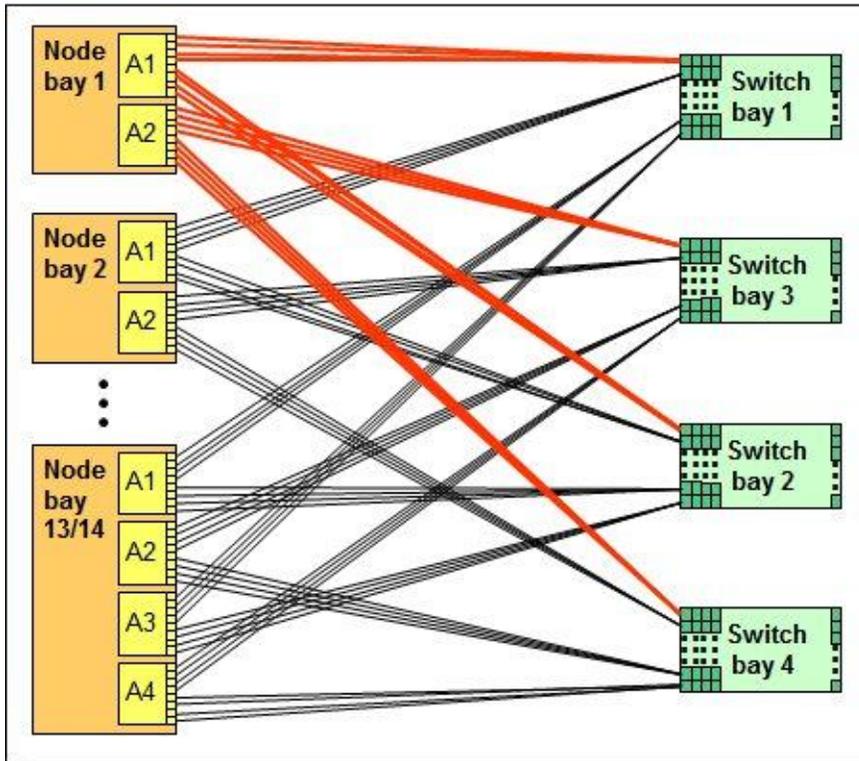


Figure 3. Logical layout of the interconnects between I/O adapters and I/O modules

The EN4093 and EN4093R switches 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 (90Y3562) 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, and 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 such a 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 5. 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 of the EN4093 or EN4093R 10 Gb switch.

** Ports 5 and 6 require Upgrade 2 of the EN4093 or EN4093R 10 Gb switch.

Ports 7 and 8 are reserved for future use.

The following table lists the I/O adapters supported by the EN4093 and EN4093R 10 GbE switches.

Table 6. Supported network adapters

Description	Part number	Feature code (x-config / e-config)	EN4093	EN4093R
10 Gb Ethernet				
Embedded 10Gb Virtual Fabric Adapter (2-port)	None	None / None	Yes*	Yes*
IBM Flex System Embedded 10Gb Virtual Fabric Upgrade (Feature on Demand to provide FCoE and iSCSI support)	90Y9310	A2TD / None	Yes	Yes
IBM Flex System CN4054 10Gb Virtual Fabric Adapter (4-port)	90Y3554	A1R1 / None	Yes	Yes
IBM Flex System CN4054 Virtual Fabric Adapter (SW Upgrade) (Feature on Demand to provide FCoE and iSCSI support)	90Y3558	A1R0 / None	Yes	Yes
IBM Flex System CN4058 8-port 10Gb Converged Adapter	None	None / EC24	Yes	Yes
IBM Flex System EN4054 4-port 10Gb Ethernet Adapter	None	None / 1762	Yes	Yes
IBM Flex System EN4132 2-port 10Gb Ethernet Adapter	90Y3466	A1QY / None	Yes	Yes
IBM Flex System EN4132 2-port 10Gb RoCE Adapter	None	None / EC26	Yes	Yes
1 Gb Ethernet				
Embedded 1 Gb Ethernet controller (2-port)	None	None / None	Yes	Yes
IBM Flex System EN2024 4-port 1Gb Ethernet Adapter	49Y7900	A10Y / 1763	Yes	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.

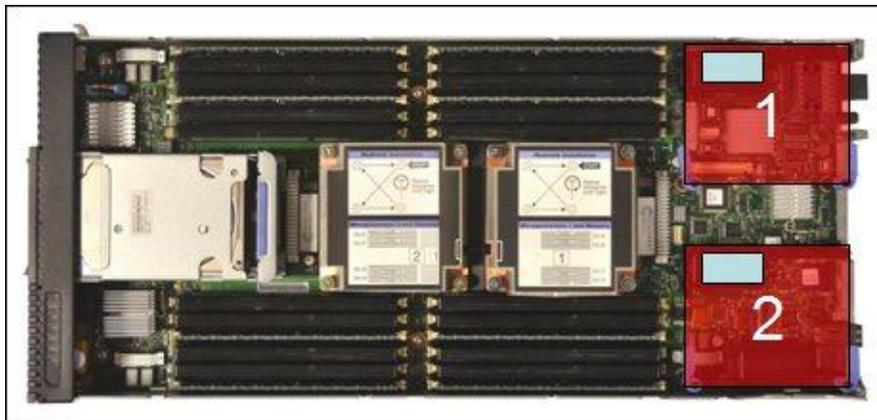


Figure 4. Location of the I/O adapter slots in the IBM Flex System x240 Compute Node

Connectors and LEDs

Figure 5 shows the front panel of the IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches.

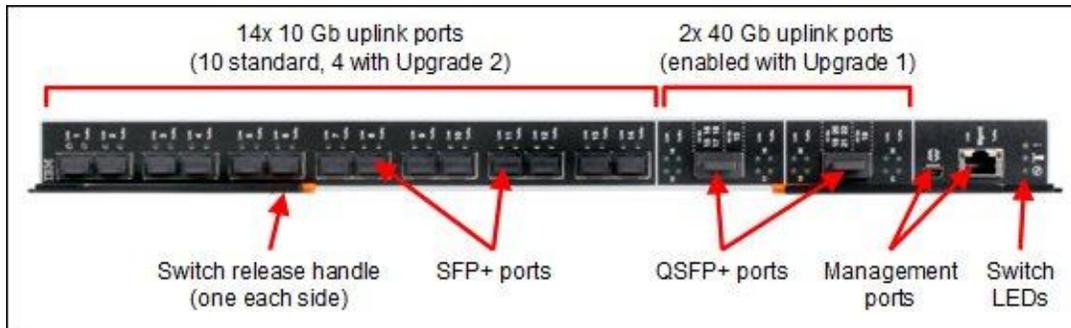


Figure 5. Front panel of the IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches

The front panel contains the following components:

- 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.)
- Fourteen external SFP+ ports for 1 Gb or 10 Gb connections to external Ethernet 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.
- An Ethernet link OK LED and an Ethernet Tx/Rx LED for each external port on the switch module.

Network cabling requirements

The network cables that can be used with the EN4093 and EN4093R switches are shown in the following table.

Table 7. EN4093 and EN4093R 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
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, these switches assume 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 Communiqué 2004/9; Communiqué 2004/22
- Saudi Arabia EMC.CVG, 28 October 2002

Popular configurations

The following usage scenarios are described:

- EN4093R as a 10 Gb Ethernet Virtual Fabric switch
- EN4093R as an FCoE transit switch

EN4093R as a 10 Gb Ethernet Virtual Fabric switch

The EN4093R 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 EN4093R 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 each vNIC. Up to 32 vNICs can be configured on a half-wide compute node.

The EN4093R switches can be connected to the top-of-rack aggregator switches:

- IBM RackSwitch G8264 via the 10 GbE uplinks
- IBM RackSwitch G8316 via the 40 GbE uplinks

The following figure illustrates this scenario.

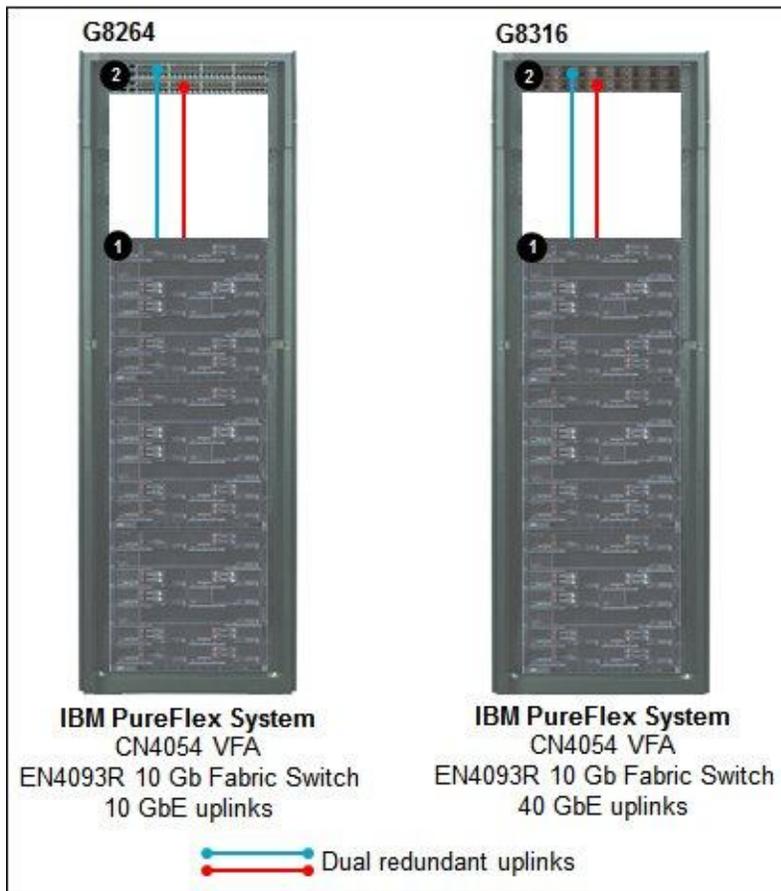


Figure 6. EN4093R as a 10 Gb Virtual Fabric Switch

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

Table 8. Components used in a Virtual Fabric solution with the EN4093R switch (Figure 6)

Diagram reference	Description	Part number	Quantity
①	IBM Flex System Virtual Fabric solution		
	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 EN4093R 10Gb Scalable Switch	95Y3309	2 per chassis
	IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	49Y4798	1 per EN4093R
②	IBM RackSwitch G8264 or G8316		

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

EN4093R as an FCoE transit switch

IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches are Data Center Bridging (DCB) switches that can transport FCoE frames by using FCoE Initialization Protocol (FIP) snooping. These switches provide an inexpensive solution for transporting encapsulated FCoE packet to the Fibre Channel Forwarder (FCF) which is functioning as both an aggregation switch and an FCoE gateway, Examples of this scenario are depicted in Figure 7 and Figure 8..

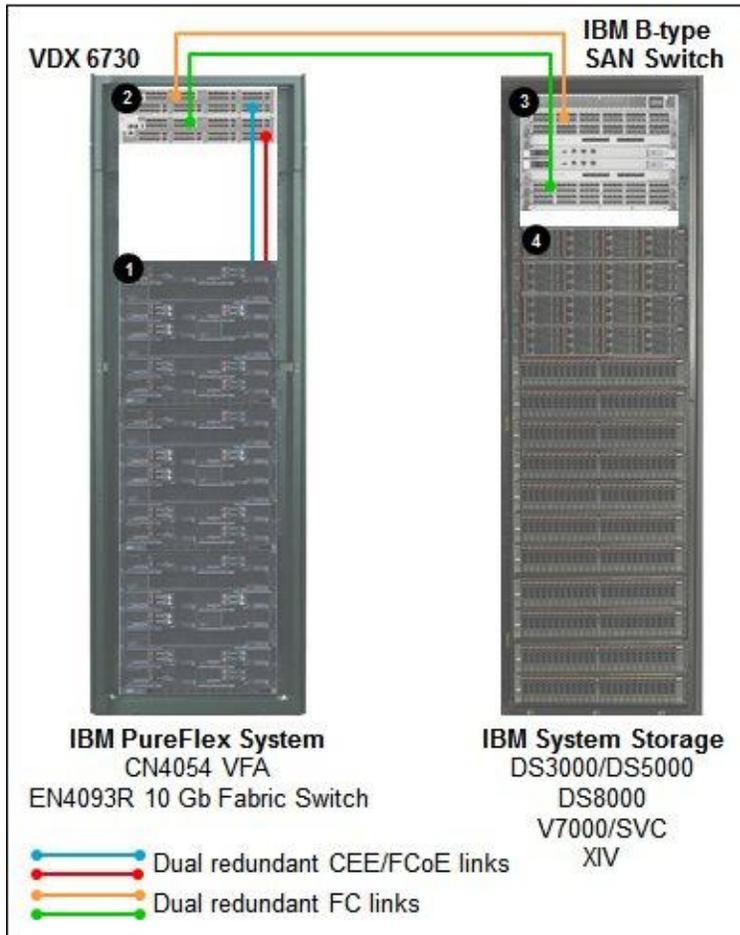


Figure 7. EN4093R as an FCoE transit switch with the Brocade VDX 6730 as an FCF

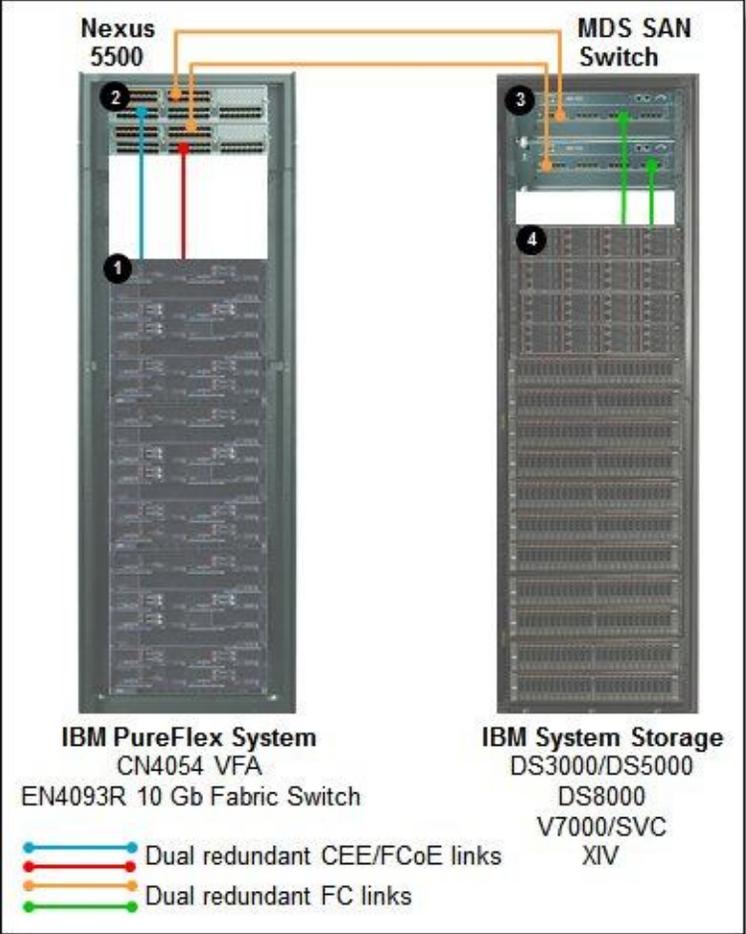


Figure 8. EN4093R as an FCoE transit switch with the Cisco Nexus 5548/5596 as an FCF

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

Table 9. EN4093R as an FCoE transit switch with the Brocade VDX 6730 as an FCF (Figure 7)

Diagram reference	Description	Part number	Quantity
1	IBM Flex System FCoE solution		
	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 EN4093R 10Gb Scalable Switch	95Y3309	2 per chassis
	IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	49Y4798	1 per EN4093R
2	Brocade VDX 6730 Converged Switch for IBM		
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 10. EN4093R as an FCoE transit switch with the Cisco Nexus 5548/5596 as an FCF (Figure 8)

Diagram reference	Description	Part number	Quantity
1	IBM Flex System FCoE solution		
	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 EN4093R 10Gb Scalable Switch	95Y3309	2 per chassis
	IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	49Y4798	1 per EN4093R
2	Cisco Nexus 5548/5596 Switch		
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 9 and Table 10; see Table 3 for details) for the external 10 Gb Ethernet 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 EN4093 and EN4093R 10Gb Scalable Switches product publications, available from the IBM Flex System Information Center at
<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 EN4093
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS112-053>
- IBM US Announcement Letter for the EN4093R
<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/>

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.

© Copyright International Business Machines Corporation 2012. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on February 5, 2013.

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/tips0864.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®
IBM Flex System™
IBM Power Systems™
Redbooks®
Redbooks (logo)®
ServicePac®
VMready®

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