Redpaper

Ilya Krutov

The Benefits of 16 Gb Fibre Channel in IBM Flex System Solutions

Ensuring that business-critical data is available when needed is an ever-growing need in IT. Your systems must store massive amounts of data quickly and retrieve it efficiently. Simultaneously, you must use new technologies that can improve efficiency and take advantage of these technologies within limited budgets.

One measure of growing efficiency in recent years is CPU processing power, which far exceeds growth in disk input/output (I/O). For this reason, disk I/O is often the reason for bottlenecks in high-performance applications.

16 Gb Fibre Channel connectivity that is combined with other storage technologies, such as solid-state drives and storage tiering, can help to balance the performance of storage subsystem with the server processing capabilities.

This IBM® Redpaper™ publication describes server performance imbalance that can be found in typical application environments and how to address this issue with the 16 Gb Fibre Channel technology to provide required levels of performance and availability for the storage-intensive applications.

The following topics are covered:

Executive summary

Currently, the processor, memory, and I/O subsystem are balanced and not considered as performance bottlenecks in most of systems. The major source of performance issues is related to the storage I/O activity because of the speed of traditional storage systems that still does not match the processing capabilities of the servers. This disparity can lead to a situation where a powerful processor sits idle waiting for the storage I/O requests to complete. This situation wastes the processor's time, which negatively impacts user productivity, extends return on investments (ROI) time frame, and increases overall total cost of ownership (TCO).

With the virtualization trends in data centers, servers demand higher I/O bandwidth to match the capabilities of multi-core processors and increased amounts of memory, allowing the higher number of virtual machines (VMs) to be hosted on a single physical system. Higher I/O bandwidth, including storage I/O, can help to achieve better server utilization and higher VM per server ratio.

Data warehouses and business analytics are additional examples of the workload that requires higher storage I/O bandwidth to allow faster data processing, making strategic business decisions in a timely manner.

IBM Flex System™ 16 Gb Fibre Channel solution can help to address these issues by increasing the speed of storage fabric by up to 1.6 GBps per port per direction. Combined with the reliable high-speed solid-state drive technology and storage tiering, 16 Gb FC fabric can help to significantly decrease storage I/O response time to match the processing power of the server CPUs.

IBM Flex System integrated 16 Gb FC technology can help to achieve:

- Up to 50 - 100% better virtual machine density and a higher number of concurrent users because of increased storage bandwidth
- Up to 20 - 40% fewer servers are required to support the workload that is specified
- Twice lower number and higher speed of inter-switch links required in scalable SANs
- Higher reliability and availability of services due to fewer number of components that are used to build the solution
- Twice faster access to the business critical data
- Lower acquisition costs due to fewer number of systems and components
- Shorten ROI time frame and decrease overall TCO with the efficient utilization of server resources and lower power, cooling, and management costs

Introduction

Storage is no longer an afterthought. Companies are searching for more ways to efficiently manage expanding volumes of data and to make that data accessible throughout the enterprise. In addition, strategic workloads, such as data analytics, which support business decisions and rely on efficient data mining, including data access speed and processing time. This demand is propelling the move of storage into the network.

With current estimates of the amount of data to be managed and made available increasing significantly each year, this outlook is where a *storage area network (SAN)* enters the arena. SANs are the leading storage infrastructure for the global economy of today. SANs offer simplified storage management, scalability, flexibility, and availability and improved data access, movement, and backup.

Storage area network

Storage area network (SAN) is a network whose primary purpose is the transfer of data between computer systems and storage elements. A SAN consists of a communication infrastructure, which provides physical connections. It also includes a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. The term SAN is identified with block I/O services rather than file access services.

Figure 1 shows an overview of a SAN that connects multiple servers to the storage systems.

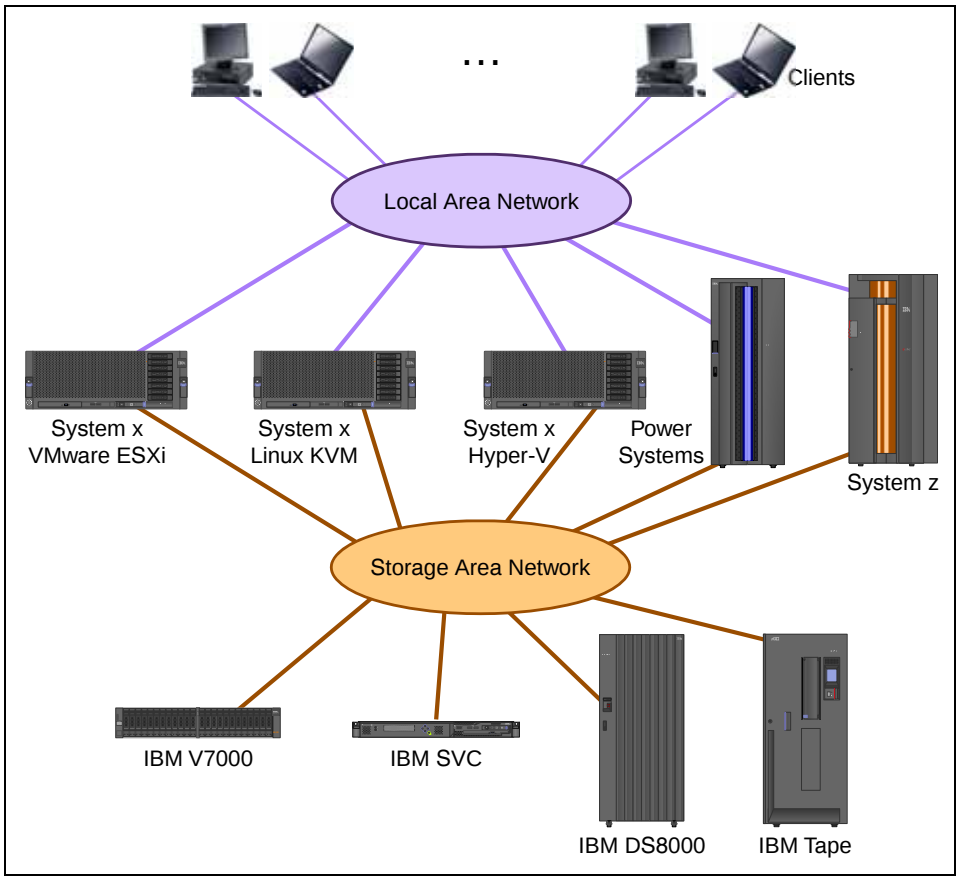

Figure 1 A storage area network

SANs provide advanced methods of attaching storage to servers. These advanced methods can enable great improvements in both availability and performance. The SANs are used to connect shared storage arrays and tape libraries to multiple servers. They are also used by clustered servers for failover.

A SAN can be used to bypass traditional network bottlenecks. It facilitates direct, high-speed data transfers between servers and storage devices in the following ways:

- Server to storage: This is the traditional model of interaction with storage devices. The advantage is that the same storage device might be accessed serially or concurrently by multiple servers.
- Storage to storage: This outboard data movement capability enables data to be moved without server intervention, freeing up server processor cycles for other activities, such as application processing. Examples include a disk device that backs up its data to a tape device without server intervention or a remote device mirroring across the SAN.

SANs allow applications that move data to perform better; for example, by having the data sent directly from the source to the target device with minimal server intervention.

SANs also enable new network architectures where multiple hosts access multiple storage devices that are connected to the same network. Using a SAN can provide the following benefits:

- Improvements to application availability: Storage is independent of applications and accessible through multiple data paths for better reliability, availability, and serviceability.
- Higher application performance: Storage processing is offloaded from servers and moved onto a separate network.
- Centralized and consolidated storage: Simpler management, scalability, flexibility, and availability.
- Data transfer and vaulting to remote sites: Remote copy of data that is enabled for disaster protection and against malicious attacks.
- Simplified centralized management: Single image of storage media simplifies management.

The storage infrastructure is the foundation on which information relies, and therefore must support the business objectives and business model of a company. In this environment, simply deploying more and faster storage devices is not enough. A SAN infrastructure provides enhanced network availability, data accessibility, and system manageability.

Storage requirements

Choosing the correct storage for application data can be a complex task because you must ensure that critical business and application requirements are met while costs are kept optimized. In particular, storage performance capabilities must match the processing capabilities of the server itself to ensure the most efficient utilization of system resources. There is no “one size fits all” approach possible because different applications have different storage data access patterns.

In general, the factors to consider during the planning process for application data storage include:

- Importance of data (Can I accept the loss of data?)
- Sensitivity of data (Do I need an advanced data protection and security?)
- Availability of data (Do I need the data 24 hours per day, 7 days per week?)
- Security of data (Who can read, modify, and delete the data?)
- Data access speed (How quickly do I need to insert and extract the data?)
- Performance or workload capacity (How many IOPS for I/O-intensive workloads and how many MBps for throughput-intensive workloads do I need?)

- Storage capacity (How much space do I need to store the data?)
- Frequency of access (How often do I need the data?)
- Backup and recovery strategy (How much time do I need to back up and restore the data?)
- Retention policy (How long do I keep the data?)
- Scalability for future growth (Do I expect the workload increase in the near future?)
- Storage deployment: internal or external (If external, then JBOD or storage controller? If storage controller, then SAS, iSCSI, FC, or FCoE?)
- Data access pattern (How does the application access the data?):
  - Read or write intensive
  - Random or sequential access
  - Large or small I/O requests

Answers to these questions help you to formalize the performance, availability, and capacity requirements for your applications and match these requirements with the appropriate storage design model.

Storage performance

In general, there are two key types of storage applications based on workloads that they generate:

- *I/O-intensive* applications require storage systems to process as many hosts read and write requests (or I/O requests) per second as possible given the average I/O request size used by this application, which is typically 8 - 16 KBytes. This behavior is most common for OLTP databases.
- *Throughput-intensive* applications require storage systems to transfer to or from host as many gigabytes of information per second as possible, and they typically use I/O request sizes of 64 - 128 KB. These characteristics are commonly inherent to file servers, multimedia streaming, and backup.

Therefore, there are two key performance metrics to evaluate storage system performance: *input/output requests per second (IOPS)* and *throughput (measured in GBps)* depending on application workload.

Table 1 lists typical storage-intensive applications and their workload patterns in a multi-user environment.

Table 1 Typical application workload patterns

Workload Type → ↓ Application Type	**Read intensive**	**Write intensive**	**I/O intensive**	**Throughput intensive**	**Random access**	**Sequential access**	**Good for 16 Gb FC**
File server	Yes			Yes	Yes		Yes
OLTP Database	Yes	Yes	Yes		Yes		
Data warehouse	Yes			Yes	Yes		Yes
Email server	Yes	Yes	Yes		Yes		
Medical imaging	Yes			Yes	Yes		Yes
Video on demand	Yes			Yes	Yes		Yes
Streaming media	Yes			Yes	Yes		Yes

Workload Type → ↓ Application Type	Read intensive	Write intensive	I/O intensive	Throughput intensive	Random access	Sequential access	Good for 16 Gb FC
Web/Internet	Yes		Yes		Yes		
Web 2.0	Yes	Yes	Yes		Yes		

Throughput-intensive applications, such as data warehouses, video on demand, and others, can benefit from 16 Gb Fibre Channel.

Fibre Channel SANs versus converged networks

In the recent past, converged networks were a popular discussion topic because of their ability to carry both LAN and SAN traffic over the same physical infrastructure. Depending on infrastructure requirements, this can help to reduce the number of ports, adapters, and devices that are used, and therefore decrease overall TCO.

However, there is no “one size fits all” approach, and the selection of storage connectivity approach can be a difficult task. Both dedicated SAN and converged networks have their own strengths and benefits, and the final choice depends on what are the most important requirements in each particular deployment.

Table 2 provides an overview of factors that can be considered when deciding on which approach to choose.

Tip: The terms Moderate and High that is used in Table 2 are relative indicators for comparison purposes and do not represent any meaning in terms of absolute values. For example, values in the TCO row mean that FC SAN typically (but not necessarily always, it depends on specific deployment requirements) has a higher TCO than converged FCoE network.

Table 2 FC SAN versus converged FCoE network

Factor	FC SAN	Converged FCoE network
Storage I/O workload	Medium to heavy	Light to medium
Storage I/O infrastructure	Dedicated	Shared
Traffic isolation	Physical	Logical
Availability	High	Moderate
Performance	High	Moderate
Scalability	High	Moderate
Security	High	Moderate
Storage infrastructure management	Separate	Unified
Acquisition costs	High	Moderate
Management costs	High	Moderate
Operational costs	High	Moderate
TCO	High	Moderate

In general, as shown in Table 2 on page 6, FC SANs provide better performance, availability, scalability, and security, and converged networks can achieve lower TCO.

If the storage workload is light to medium and TCO is the key decision factor, a converged network is an appropriate choice. If the storage workload is moderate to heavy and performance, availability, scalability, and security are the key decision factors, FC SAN is an appropriate choice.

Benefits of 16 Gb Fibre Channel SANs

In general, 16 Gb Fibre Channel means twice higher bandwidth compared to the 8 Gb Fibre Channel. As was already described, the processor, memory, and I/O subsystem are balanced and not considered as performance bottlenecks in the majority of systems. The main source of performance issues tends to be related to storage I/O activity because the speed of traditional storage systems still does not match the processing capabilities of the servers. Therefore, increasing the speed of storage connections can help to close or minimize this performance gap.

The use of 16 Gb FC connectivity can help to achieve the following benefits:

- Higher performance:
  - Higher IOPS
  - Higher throughput
  - Lower latency
- Infrastructure simplification:
  - Simplified deployment and management
  - Higher VM density/lower number of physical systems
- Improved TCO:
  - Reduced acquisition costs
  - Reduced power and cooling costs
  - Reduced support/maintenance costs

Performance

The maximum theoretical throughput of 16 Gb Fibre Channel is 1.6 GBps, and 8 Gb FC has a throughput of 0.8 GBps. This fact means that throughput-intensive applications can achieve better response time (for example, with OLAP workloads) or support larger numbers of concurrent users on a single server (such as Video on Demand or streaming media workloads).

Also, 16 Gb FC can potentially be beneficial for IOPS-intensive applications: 16 Gb infrastructure can deliver up to 200,000 IOPS using the 8 KB I/O blocks. 8 Gb can deliver only up to 100,000 IOPS.

Infrastructure simplification

Because the performance of a single 16 Gb FC port is equivalent to the performance of two 8 Gb FC ports, you can reduce the number of ports, cables, and SFP modules while keeping the same performance levels.

Also, if 8 Gb FC is a limiting factor for overall server performance, you can potentially support heavier workloads on a same system by moving to 16 Gb, increasing VM density or concurrent user support while lowering overall number of physical systems.

Improved TCO

With the lower number of systems and components that are used to build the infrastructure, you can potentially reduce acquisition costs and operational, management, and support costs due to simplified infrastructure.

This paper describes several possible deployment scenarios where IBM Flex System 16 Gb Fibre Channel offerings can provide certain benefits:

- Data warehouse
- Video on Demand
- Mixed virtualized workload
- Dual VIOS configuration
- Backup solutions with separated disk and tape traffic
- x222 Fibre Channel storage connectivity

For more information, see “Deployment scenarios of 16 Gb FC host connectivity” on page 25.

IBM Flex System I/O architecture and 16 Gb FC portfolio

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 costs per workload. Although increased security and resiliency protect vital information and promote maximum uptime, the integrated, easy-to-use management system reduces set up time and complexity, providing a quicker path to return on investment (ROI).

The following topics are described in this section:

- Enterprise Chassis I/O architecture
- IBM Flex System 16 Gb FC components

Enterprise Chassis I/O architecture

The IBM Flex System Enterprise Chassis is a 10U next-generation server platform with integrated chassis management. It is a compact, high-density, high-performance, rack-mount, scalable platform system. It supports up to 14 one-bay compute nodes that share common resources, such as power, cooling, management, and I/O resources within a single Enterprise Chassis. In addition, it can also support up to seven 2-bay compute nodes or three 4-bay compute nodes when the shelves are removed. You can mix and match 1-bay, 2-bay, and 4-bay compute nodes to meet your specific hardware needs.

From a physical I/O module bay perspective, the Enterprise Chassis has four I/O bays in the rear of the chassis. Figure 2 shows the physical layout of these I/O module bays.

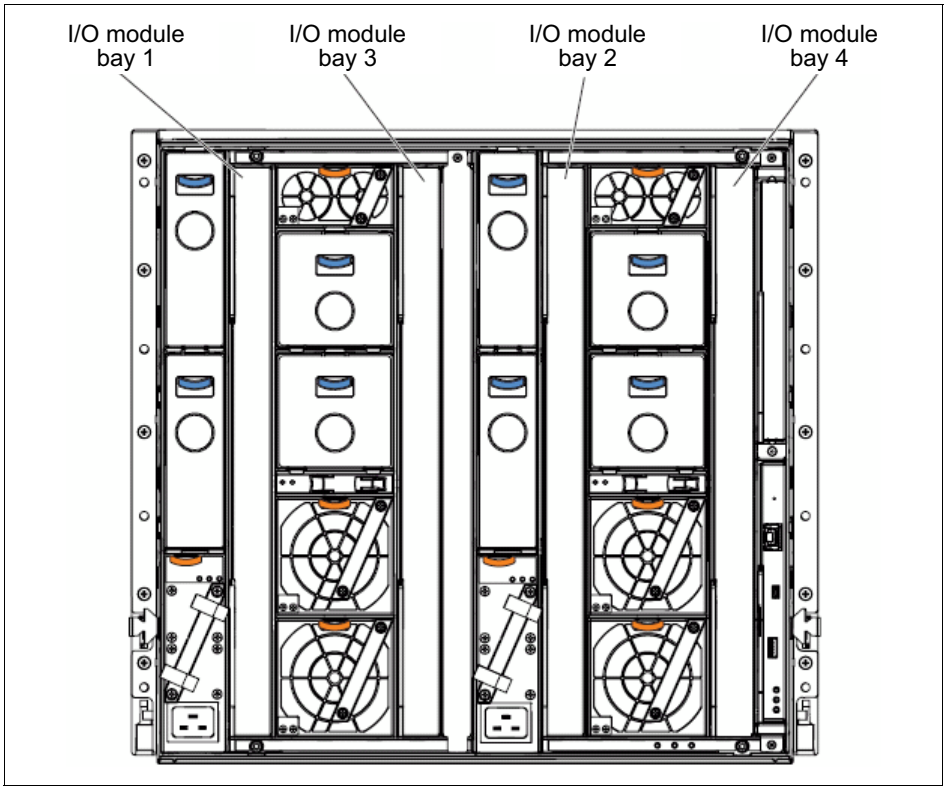

Figure 2 Rear view of the Enterprise Chassis showing I/O module bays

If a node has a two-port integrated LAN on Motherboard (LOM) as standard, Module 1 and 2 are connected to this LOM. If an I/O adapter is installed in the node's I/O expansion slot 1, Modules 1 and 2 are connected to this adapter.

Modules 3 and 4 connect to the I/O adapter that is installed within I/O expansion bay 2 on the node.

These I/O modules provide external connectivity and connect internally to each of the nodes within the chassis.

The connections between the 2-port adapters that are installed in the compute nodes and the I/O bays in the chassis are shown in Figure 3. Figure 3 also shows both half-wide servers, such as the x240 with two adapters, and full-wide servers, such as the x440 with four adapters.

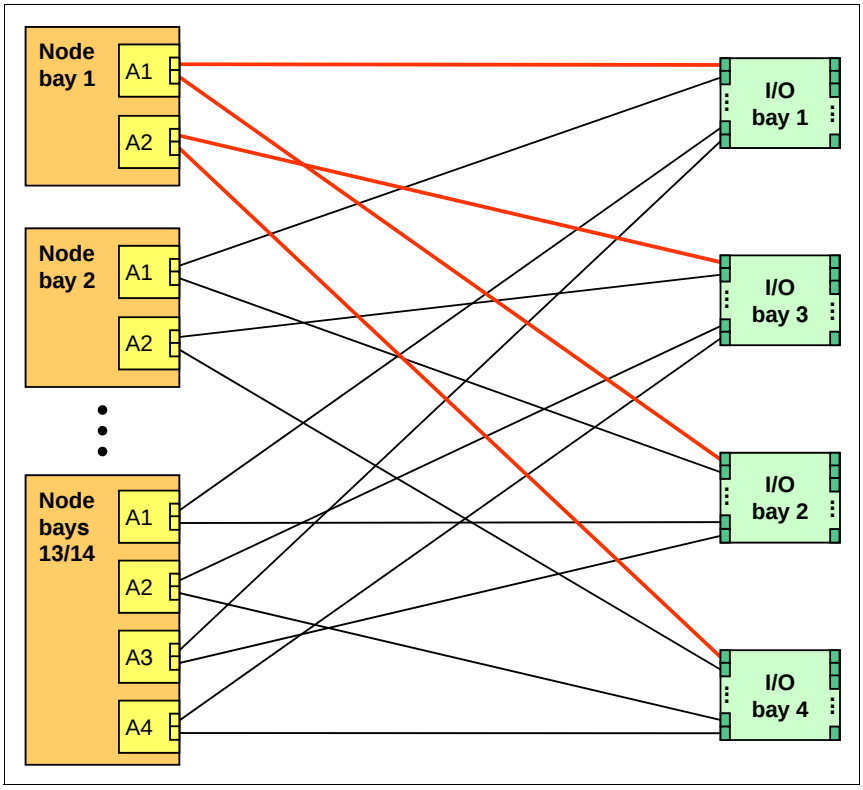

Figure 3 Logical layout of node to switch interconnects for 2-port cards

A total of two I/O expansion adapters (designated A1 and A2 in Figure 3) can be plugged into a half-wide node. Up to four I/O adapters can be plugged into a full-wide node.

Figure 4 shows the links from the 4-port I/O adapters to the I/O modules.

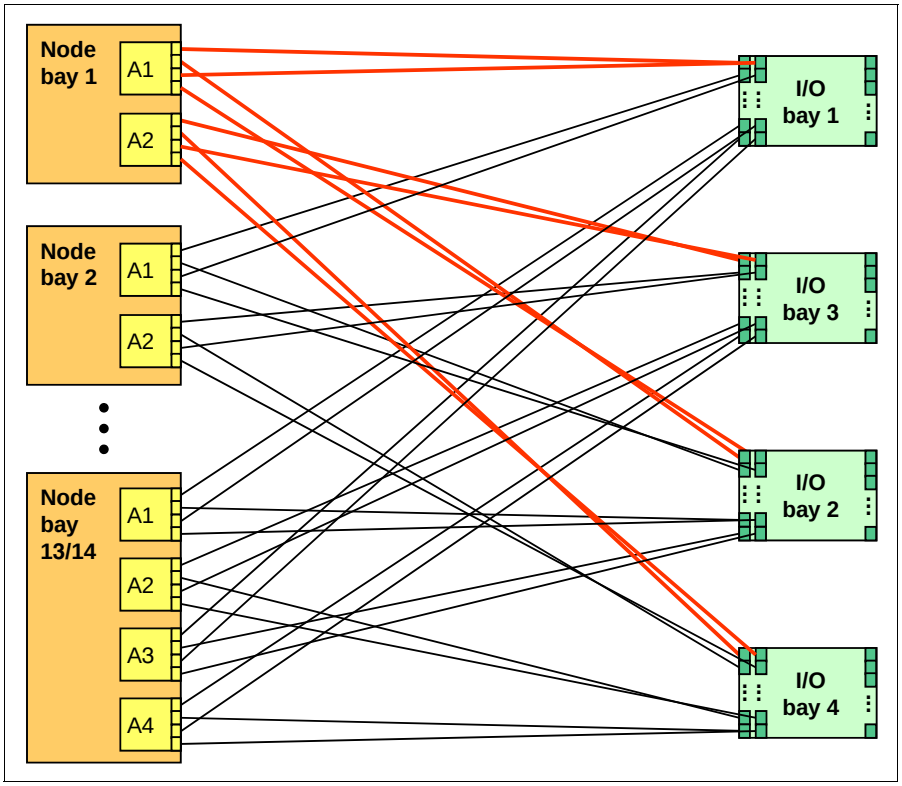

Figure 4 Logical layout of node to switch interconnects for 4-port cards

Each I/O adapter has two connectors. One connector attaches to the compute node's system board (a PCI Express connection). The second connector is a high-speed interface to the midplane that mates to the midplane when the node is installed into a bay within the chassis.

Fibre Channel components of IBM Flex System are as follows:

- FC adapters are typically installed in slot #2 of the half-wide compute node
- FC switches are typically installed in I/O bays 3 and 4

IBM Flex System 16 Gb FC components

This section describes the following 16 Gb FC components that are available for deployment with the Enterprise Chassis:

- IBM Flex System FC5022 16 Gb SAN Scalable Switch
- 16 Gb FC adapters

IBM Flex System FC5022 16 Gb SAN Scalable Switch

The IBM Flex System FC5022 16 Gb SAN Scalable Switch is a high-density, 48-port, 16 Gbps Fibre Channel switch that is used in the Enterprise Chassis. The switch provides 28 internal ports to compute nodes by way of the midplane and 20 external SFP+ ports. These system area network (SAN) switch modules deliver an embedded option for IBM Flex System users who deploy storage area networks in their enterprise. They offer end-to-end 16 Gb and 8 Gb connectivity.

The N_Port Virtualization mode streamlines the infrastructure by reducing the number of domains to manage. It allows you to add or move servers without impact to the SAN.

Monitoring is simplified by using an integrated management appliance. Clients who use end-to-end Brocade SAN can take advantage of the Brocade management tools.

Figure 5 shows the IBM Flex System FC5022 16 Gb SAN Scalable Switch.

Figure 5 IBM Flex System FC5022 16 Gb SAN Scalable Switch

Three versions are available as listed in Table 3: 12-port and 24-port switch modules and a 24-port switch with the Enterprise Switch Bundle (ESB) software. The port count can be applied to internal or external ports by using a feature that is called Dynamic Ports on Demand (DPOD). Ports counts can be increased with license upgrades as described in “Port and feature upgrades” on page 13.

Table 3 IBM Flex System FC5022 16 Gb SAN Scalable Switch part numbers

Part number	Feature codes[a]	Description	Ports enabled by default
88Y6374	A1EH / 3770	IBM Flex System FC5022 16 Gb SAN Scalable Switch	12
00Y3324	A3DP / ESW5	IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switch	24
90Y9356	A1EJ / 3771	IBM Flex System FC5022 24-port 16 Gb ESB SAN Scalable Switch	24

a. x-config (HVEC) / e-config (AAS) feature code.

Table 4 provides a feature comparison between the FC5022 switch models.

Table 4 Feature comparison by model

Feature	FC5022 16 Gb 24-port ESB Switch	FC5022 24-port 16 Gb SAN Scalable Switch	FC5022 16 Gb SAN Scalable Switch
	90Y9356	00Y3324	88Y6374
Number of active ports	24	24	12
Number of SFP+ included	None	2x 16 Gb SFP+	None
12-port Upgrade	Included	Included	Optional
24-port Upgrade	Optional	Optional	Optional
Full fabric	Included	Included	Included
Access Gateway	Included	Included	Included
Advanced zoning	Included	Included	Included
Enhanced Group Management	Included	Included	Included
ISL Trunking	Included	Optional	Optional
Adaptive Networking	Included	Not available	Not available

Feature	FC5022 16 Gb 24-port ESB Switch	FC5022 24-port 16 Gb SAN Scalable Switch	FC5022 16 Gb SAN Scalable Switch
	90Y9356	00Y3324	88Y6374
Advanced Performance Monitoring	Included	Not available	Not available
Fabric Watch	Included	Optional	Optional
Extended Fabrics	Included	Not available	Not available
Server Application Optimization	Included	Not available	Not available

The part number for the switch includes the following items:

- One IBM Flex System FC5022 16 Gb SAN Scalable Switch or IBM Flex System FC5022 24-port 16 Gb ESB SAN Scalable Switch
- Important Notices Flyer
- Warranty Flyer
- Documentation CD-ROM

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 cable can be used to connect to the switch locally for configuration tasks and firmware updates.

Port and feature upgrades

Table 5 lists the available port and feature upgrades. These are all IBM Features on Demand license upgrades.

Table 5 FC5022 switch upgrades

Part number	Feature codes[a]	Description	24-port 16 Gb ESB switch	24-port 16 Gb SAN switch	16 Gb SAN switch
			90Y9356	00Y3324	88Y6374
88Y6382	A1EP / 3772	FC5022 16 Gb SAN Switch (Upgrade 1)	No	No	Yes
88Y6386	A1EQ / 3773	FC5022 16 Gb SAN Switch (Upgrade 2)	Yes	Yes	Yes
00Y3320	A3HN / ESW3	FC5022 16 Gb Fabric Watch Upgrade	No	Yes	Yes
00Y3322	A3HP / ESW4	FC5022 16 Gb ISL/Trunking Upgrade	No	Yes	Yes

a. x-config (HVEC) / e-config (AAS) feature code.

With DPOD, ports are licensed as they come online. With the FC5022 16 Gb SAN Scalable Switch, the first 12 ports that report (on a first-come, first-served basis) on boot-up are assigned licenses. These 12 ports can be any combination of external or internal Fibre Channel ports. After all licenses are assigned, you can manually move those licenses from one port to another. Because this process is dynamic, no defined ports are reserved except ports 0 and 29. The FC5022 16Gb ESB Switch has the same behavior. The only difference is the number of ports.

Table 6 shows the total number of active ports on the switch after applying compatible port upgrades.

Table 6 Total port counts after applying upgrades

	Total number of active ports		
	24-port 16 Gb ESB SAN switch	**24-port 16 Gb SAN switch**	**16 Gb SAN switch**
Ports on Demand upgrade	**90Y9356**	**00Y3324**	**88Y6374**
Included with base switch	24	24	12
Upgrade 1, 88Y6382 (adds 12 ports)	Not supported (included)	Not supported (included)	24
Upgrade 2, 88Y6386 (adds 24 ports)	48	48	48

Transceivers

The FC5022 12-port and 24-port ESB SAN switches come without SFP+, which must be ordered separately to provide outside connectivity. The FC5022 24-port SAN switch comes standard with two Brocade 16 Gb SFP+ transceivers, and additional SFP+ can be ordered if required. Table 7 lists supported SFP+ options.

Table 7 Supported SFP+ transceivers

Part number	Feature code[a]	Description
16 Gb SFP+ transceivers		
88Y6393	A22R / 5371	Brocade 16 Gb SFP+ Optical Transceiver
98Y2178	None / 2611	SFP+ Transceiver 16 Gbps 10 km LW
98Y2179	None / 2618	SFP+ Transceiver 16 Gbps 10 km LW 8-Pack
8 Gb SFP+ transceivers		
88Y6416	5084 / 5370	Brocade 8 Gb SFP+ SW Optical Transceiver
45W1216	None / 2821	SFP Transceiver 8 Gbps 10 km LW
45W1218	None / 2828	SFP Transceiver 8 Gbps 10 km LW 8-Pack
45W2283	None / 2881	SFP Transceiver 8 Gbps 25 km ELW

a. x-config (HVEC) / e-config (AAS) feature code.

Benefits

The switches offer the following key benefits:

- Exceptional price/performance for growing SAN workloads

  The FC5022 16 Gb SAN Scalable Switch delivers exceptional price/performance for growing SAN workloads. It achieves this through a combination of market-leading 1.6 GBps throughput per port and an affordable high-density form factor. The 48 FC ports produce an aggregate 768 Gbps full-duplex throughput, plus any external eight ports can be trunked for 128 Gbps inter-switch links (ISLs). Because 16 Gbps port technology dramatically reduces the number of ports and associated optics/cabling required through 8/4 Gbps consolidation, the cost savings and simplification benefits are substantial.

- Accelerating fabric deployment and serviceability with diagnostic ports

  Diagnostic Ports (D_Ports) are a new port type that is supported by the FC5022 16 Gb SAN Scalable Switch. They enable administrators to quickly identify and isolate 16 Gbps optics, port, and cable problems, reducing fabric deployment and diagnostic times. If the optical media is found to be the source of the problem, it can be transparently replaced because 16 Gbps optics are hot-pluggable.

- A building block for virtualized, private cloud storage

  The FC5022 16 Gb SAN Scalable Switch supports multi-tenancy in cloud environments through VM-aware end-to-end visibility and monitoring, QoS, and fabric-based advanced zoning features. The FC5022 16 Gb SAN Scalable Switch enables secure distance extension to virtual private or hybrid clouds with dark fiber support. They also enable in-flight encryption and data compression. Internal fault-tolerant and enterprise-class reliability, availability, and serviceability (RAS) features help minimize downtime to support mission-critical cloud environments.

- Simplified and optimized interconnect with Brocade Access Gateway

  The FC5022 16 Gb SAN Scalable Switch can be deployed as a full-fabric switch or as a Brocade Access Gateway. It simplifies fabric topologies and heterogeneous fabric connectivity. Access Gateway mode uses N_Port ID Virtualization (NPIV) switch standards to present physical and virtual servers directly to the core of SAN fabrics. This configuration makes it not apparent to the SAN fabric, greatly reducing management of the network edge.

- Maximizing investments

  To help optimize technology investments, IBM offers a single point of serviceability that is backed by industry-renowned education, support, and training. In addition, the IBM 16/8 Gbps SAN Scalable Switch is in the IBM ServerProven® program, enabling compatibility among various IBM and IBM Business Partner products. IBM recognizes that customers deserve the most innovative, expert integrated systems solutions.

Features and specifications

FC5022 16 Gb SAN Scalable Switches have the following features and specifications:

- Internal ports:
  - 28 internal full-duplex 16 Gb FC ports (up to 14 internal ports can be activated with Port-on-Demand feature, remaining ports are reserved for future use)
  - Internal ports operate as F_ports (fabric ports) in native mode or in access gateway mode
  - Two internal full-duplex 1 GbE ports connect to the chassis management module
- External ports:
  - Twenty external ports for 16 Gb SFP+ or 8 Gb SFP+ transceivers that support 4 Gb, 8 Gb, and 16 Gb port speeds. SFP+ modules are not included and must be purchased separately. Ports are activated with the Port-on-Demand feature.
  - External ports can operate as F_ports, FL_ports (fabric loop ports), or E_ports (expansion ports) in native mode. They can operate as N_ports (node ports) in access gateway mode.
  - One external 1 GbE port (1000BASE-T) with RJ-45 connector for switch configuration and management.
  - One RS-232 serial port (mini-USB connector) that provides an additional means to configure the switch module.
- Access gateway mode (N_Port ID Virtualization - NPIV) support

- Power-on self-test diagnostics and status reporting
- ISL Trunking (licensable) allows up to eight ports (at 16, 8, or 4 Gbps speeds) to combine. These ports form a single, logical ISL with a speed of up to 128 Gbps (256 Gbps full duplex). This configuration allows for optimal bandwidth utilization, automatic path failover, and load balancing.
- Brocade Fabric OS delivers distributed intelligence throughout the network and enables a wide range of applications with added value. These applications include Brocade Advanced Web Tools and Brocade Advanced Fabric Services (on certain models).
- Supports up to 768 Gbps I/O bandwidth
- 420 million frames switch per second, 0.7 microseconds latency
- 8,192 buffers for up to 3,750 km extended distance at 4 Gbps FC (Extended Fabrics license is required to support longer distances)
- In-flight 64 Gbps Fibre Channel compression and decompression support on up to two external ports (no license required)
- In-flight 32 Gbps encryption and decryption on up to two external ports (no license is required)
- 48 Virtual Channels per port
- Port mirroring to monitor ingress or egress traffic from any port within the switch
- Two I2C connections able to interface with redundant management modules
- Hot pluggable, up to four hot pluggable switches per chassis
- Single fuse circuit
- Four temperature sensors
- Managed with Brocade Web Tools
- Supports a minimum of 128 domains in Native mode and Interoperability mode
- Nondisruptive code load in Native mode and Access Gateway mode
- 255 N_port logins per physical port
- D_port support on external ports
- Class 2 and Class 3 frames
- SNMP v1 and v3 support
- SSH v2 support
- Secure Sockets Layer (SSL) support
- NTP client support (NTP V3)
- FTP support for firmware upgrades
- SNMP/Management Information Base (MIB) monitoring functionality is contained within the Ethernet Control MIB-II (RFC1213-MIB)
- End-to-end optics and link validation
- Sends switch events and syslogs to the CMM
- Traps identify cold start, warm start, link up/link down and authentication failure events
- Support for IPv4 and IPv6 on the management ports

The FC5022 16 Gb SAN Scalable Switches come standard with the following software features:

- Brocade Full Fabric mode: Enables high performance 16 Gb or 8 Gb fabric switching
- Brocade Access Gateway mode: Uses NPIV to connect to any fabric without adding switch domains to reduce management complexity
- Dynamic Path Selection: Enables exchange-based load balancing across multiple Inter-Switch Links for superior performance
- Brocade Advanced Zoning: Segments a SAN into virtual private SANs to increase security and availability
- Brocade Enhanced Group Management: Enables centralized and simplified management of Brocade fabrics through IBM Network Advisor

Enterprise Switch Bundle software licenses

The IBM Flex System FC5022 24-port 16 Gb ESB SAN Scalable Switch includes a complete set of licensed features. These features maximize performance, ensure availability, and simplify management for the most demanding applications and expanding virtualization environments.

This switch comes with 24 port licenses that can be applied to either internal or external links on this switch.

This switch also includes the following ESB software licenses:

- Brocade Extended Fabrics

  Provides up to 1000km of switches fabric connectivity over long distances.

- Brocade ISL Trunking

  Allows you to aggregate multiple physical links into one logical link for enhanced network performance and fault tolerance.

- Brocade Advanced Performance Monitoring

  Enables performance monitoring of networked storage resources. This license includes the TopTalkers feature.

- Brocade Fabric Watch

  Monitors mission-critical switch operations. Fabric Watch now includes the new Port Fencing capabilities.

- Adaptive Networking

  Adaptive Networking provides a rich set of capabilities to the data center or virtual server environments. It ensures high-priority connections to obtain the bandwidth necessary for optimum performance, even in congested environments. It optimizes data traffic movement within the fabric by using Ingress Rate Limiting, quality of service, and Traffic Isolation Zones.

- Server Application Optimization (SAO)

  This license optimizes overall application performance for physical servers and virtual machines. SAO, when deployed with Brocade Fibre Channel host bus adapters (HBAs), extends Brocade Virtual Channel technology from fabric to the server infrastructure. This license delivers application-level, fine-grain QoS management to the HBAs and related server applications.

Supported Fibre Channel standards

The switches support the following Fibre Channel standards:

- FC-AL-2 INCITS 332: 1999
- FC-GS-5 ANSI INCITS 427, and includes:
  - FC-GS-4 ANSI INCITS 387: 2004
- FC-IFR INCITS 1745-D, revision 1.03 (under development)
- FC-SW-4 INCITS 418:2006
- FC-SW-3 INCITS 384: 2004
- FC-VI INCITS 357: 2002
- FC-TAPE INCITS TR-24: 1999
- FC-DA INCITS TR-36: 2004, and includes:
  - FC-FLA INCITS TR-20: 1998
  - FC-PLDA INCIT S TR-19: 1998
- FC-MI-2 ANSI/INCITS TR-39-2005
- FC-PI INCITS 352: 2002
- FC-PI-2 INCITS 404: 2005
- FC-PI-4 INCITS 1647-D, revision 7.1 (under development)
- FC-PI-5 INCITS 479: 2011
- FC-FS-2 ANSI/INCITS 424:2006, and includes:
  - FC-FS INCITS 373: 2003
- FC-LS INCITS 433: 2007
- FC-BB-3 INCITS 414: 2006
- FC-BB-2 INCITS 372: 2003
- FC-SB-3 INCITS 374: 2003 (replaces FC-SB ANSI X3.271: 1996 and FC-SB-2 INCITS 374: 2001)
- RFC 2625 IP and ARP Over FC
- RFC 2837 Fabric Element MIB
- MIB-FA INCITS TR-32: 2003
- FCP-2 INCITS 350: 2003 (replaces FCP ANSI X3.269: 1996)
- SNIA Storage Management Initiative Specification (SMI-S) Version 1.2 (includes the following items):
  - SNIA Storage Management Initiative Specification (SMI-S) Version 1.03 ISO standard IS24775-2006. (replaces ANSI INCITS 388: 2004)
  - SNIA Storage Management Initiative Specification (SMI-S) Version 1.1.0
  - SNIA Storage Management Initiative Specification (SMI-S) Version 1.2.0

For more information, see the IBM Redbooks Product Guide *IBM Flex System FC5022 16 Gb SAN Scalable Switch*, TIPS0870, which can be found at the following website:

http://www.redbooks.ibm.com/abstracts/tips0870.html?Open

16 Gb FC adapters

The 16 Gb adapter portfolio for IBM Flex System consists of four adapters that are described in this section. The following topics are covered:

- Adapter selection considerations
- IBM Flex System FC5022 2-port 16Gb FC Adapter
- IBM Flex System FC5024D 4-port 16Gb FC Adapter
- IBM Flex System FC5052 2-port and FC5054 4-port 16Gb FC Adapters

Adapter selection considerations

Table 8 lists common selection considerations that might be useful to you when you select an appropriate 16 Gb FC adapter.

Table 8 Adapter selection considerations

Feature	FC5022	FC5024D[a]	FC5052	FC5054
General purpose 16 Gb FC connectivity	Yes	Yes	Yes	Yes
Number of ports	2	4	2	4
Redundant storage connectivity	Yes	Yes	Yes	Yes
I/O virtualization (NPIV and SR-IOV)	Yes	Yes	Yes	Yes
Boot from SAN	Yes	Yes	Yes	Yes
IBM Fabric Manager support	Yes	Yes	Yes	Yes
Server Application Optimization (SAO)	Yes	Yes	No	No
Dual VIOS configuration	No	No	No	Yes
Separation of disk and tape traffic with full redundancy	No	No	No	Yes
Common driver model with 10 GbE LOM, CN4054	No	No	Yes	Yes

a. FC5024D is supported only in the x222 Compute Node.

IBM Flex System FC5022 2-port 16Gb FC Adapter

The network architecture on the IBM Flex System platform is designed to address network challenges. It gives you a scalable way to integrate, optimize, and automate your data center. The IBM Flex System FC5022 2-port 16Gb FC Adapter enables high-speed access to external SANs. This adapter is based on Brocade architecture, and offers end-to-end 16 Gb connectivity to SAN. It can auto-negotiate, and also work at 8 Gb and 4 Gb speeds. It also has enhanced features like N-port trunking.

Table 9 lists the ordering part number and feature code.

Table 9 IBM Flex System FC5022 2-port 16 Gb FC Adapter ordering information

Part number	HVEC feature code (x-config)	AAS feature code (e-config)[a]	Description
88Y6370	A1BP	EC2B / A1BP	IBM Flex System FC5022 2-port 16Gb FC Adapter

a. There are two e-config (AAS) feature codes for some options. The first is for the x240, p24L, p260 and p460 (when supported). The second is for the x220 and x440.

The IBM Flex System FC5022 2-port 16Gb FC Adapter has the following features:

- Over 500,000 IOPS per port, which maximizes transaction performance and density of VMs per compute node
- Achieves performance of 315,000 IOPS for email exchange and 205,000 IOPS for SQL Database
- Boots from SAN allows the automation SAN Boot LUN discovery to simplify boot from SAN and reduce image management complexity
- Brocade Server Application Optimization (SAO) provides quality of service (QoS) levels assignable to VM applications
- Direct I/O enables native (direct) I/O performance by allowing VMs to bypass the hypervisor and communicate directly with the adapter
- Brocade Network Advisor simplifies and unifies the management of the Brocade adapter, SAN, and LAN resources through a single pane-of-glass
- LUN Masking, an Initiator-based LUN masking for storage traffic isolation
- NPIV allows multiple host initiator N_Ports to share a single physical N_Port, dramatically reducing SAN hardware requirements
- Target Rate Limiting (TRL) throttles data traffic when accessing slower speed storage targets to avoid back pressure problems
- RoHS-6 compliant

Figure 6 shows the IBM Flex System FC5022 2-port 16Gb FC Adapter.

Figure 6 IBM Flex System FC5022 2-port 16Gb FC Adapter

For more information, see the IBM Redbooks® Product Guide *IBM Flex System FC5022 2-port 16Gb FC Adapter*, TIPS0891, which can be found at the following website:

http://www.redbooks.ibm.com/abstracts/tips0891.html?Open

IBM Flex System FC5024D 4-port 16Gb FC Adapter

Important: The IBM Flex System FC5024D 4-port 16Gb FC Adapter is supported only in the x222 Compute Node.

The IBM Flex System FC5024D 4-port 16Gb FC Adapter is a quad-port mid-mezzanine card for the IBM Flex System x222 Compute Node with two ports that are routed to each server in the x222. This adapter is based on the Brocade architecture, and offers end-to-end 16 Gb connectivity to a SAN. It has enhanced features such as N_Port trunking and N_Port ID Virtualization (NPIV) and boot-from-the-SAN with automatic LUN discovery and end-to-end SAO.

Table 10 lists the ordering part number and feature code.

Table 10 IBM Flex System FC5024D 4-port 16 Gb FC Adapter ordering information

Part number	Feature code[a]	Description
95Y2379	A3HU	IBM Flex System FC5024D 4-port 16Gb FC Adapter

a. The feature code listed is for both x-config (HVEC) and e-config (AAS).

The FC5024D works best with the IBM Flex System FC5022 16Gb SAN Scalable Switch. Working together, these components deliver considerable value by simplifying the deployment of server and SAN resources, reducing infrastructure and operational costs, and maximizing server and SAN reliability, availability, and resiliency.

The IBM Flex System FC5024D 4-port 16Gb FC Adapter has the following features:

- Supported in the dual-server x222 Compute Node, where two ports of the adapter are routed to each of the servers
- Dual ASIC design
- RoHS-6 compliant adapter

Each ASIC connects to one of the two servers in the x222 and act as two independent 2-port adapters, with the following features and functions:

- Based on the Brocade Catapult2 ASIC.
- Over 500,000 IOPS per port: Maximizes transaction performance and density of VMs per compute node.
- Achieves performance of 330,000 IOPS for email exchange and 205,000 IOPS for SQL Database.
- Boot from SAN allows the automation of SAN Boot LUN discovery to simplify boot from SAN and reduce image management complexity.
- Brocade SAO provides QoS levels that are assignable to VM applications.
- Direct I/O enables native (direct) I/O performance by allowing VMs to bypass the hypervisor and communicate directly with the adapter.
- Brocade Network Advisor simplifies and unifies the management of Brocade adapter, SAN, and LAN resources through a single interface.
- Provides LUN Masking, which is an Initiator-based LUN masking for storage traffic isolation.
- N_Port Id Virtualization (NPIV) allows multiple host initiator N_Ports to share a single physical N_Port, dramatically reducing SAN hardware requirements.

- Target Rate Limiting (TRL) throttles data traffic when accessing slower speed storage targets to avoid back pressure problems.
- A unified driver across all Brocade-based IBM adapter products with automated version synchronization capability.
- FEC provides a method to recover from errors that are caused on links during data transmission.
- Buffer-to-Buffer (BB) Credit Recovery enables ports to recover lost BB credits.
- FCP-IM I/O Profiling allows users to analyze traffic patterns and help fine-tune Fibre Channel adapter ports, fabrics, and targets for better performance.

Figure 7 shows the IBM Flex System FC5024D 4-port 16Gb FC Adapter.

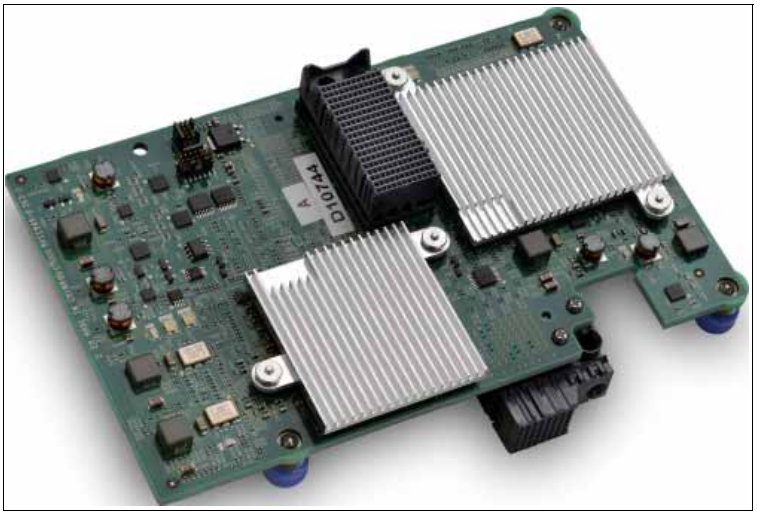

Figure 7 IBM Flex System FC5024D 4-port 16Gb FC Adapter

For more information, see *IBM Flex System FC5024D 4-port 16Gb FC Adapter*, TIPS1047, which is available at the following website:

http://www.redbooks.ibm.com/abstracts/tips1047.html?Open

IBM Flex System FC5052 2-port and FC5054 4-port 16Gb FC Adapters

The network architecture on the IBM Flex System platform addresses network challenges and give a scalable way to integrate, optimize, and automate the data center. The IBM Flex System FC5052 2-port and FC5054 4-port 16Gb FC Adapters enable high-speed access for Flex System compute nodes to an external SAN. These adapters are based on the Emulex Fibre Channel stack, and work with 16 Gb Flex System Fibre Channel switch modules.

The FC5054 adapter is based on a two ASIC design, which allows for logical partitioning on IBM Power Systems™ compute nodes. When compared to the previous generation 8 Gb adapters, the new generation 16 Gb adapters double throughput speeds for Fibre Channel traffic. As a result, it is possible to manage increased amounts of data.

You can deploy faster and manage less when you combine the FC5052 or FC5054 adapters with the CN4054 Virtual Fabric Adapter or Flex System x240 and x440 compute node's LOM, which is provided by Emulex. They use the same installation and configuration process, and Emulex OneCommand Manager enables management from a single console throughout the data center.

Table 11 lists the ordering part numbers and feature codes.

Table 11 Ordering information

Part number	x86 nodes feature[a]	IBM POWER® nodes feature	7863-10X feature	Description
95Y2386	A45R	EC23	None	FC5052 2-port 16Gb FC Adapter
95Y2391	A45S	EC2E	None	FC5054 4-port 16Gb FC Adapter

a. For all x86 compute nodes in XCC (x-config) and AAS (e-config), except for x240 7863-10X

Both adapters offer the following features:

- Fibre Channel protocol SCSI (FCP-SCSI) and Fibre Channel Internet Protocol (FCP-IP).
- Point-to-point fabric connection: F-Port Fabric Login.
- Fibre Channel Arbitrated Loop (FC-AL) and FCAL-2 FL-Port Login.
- Fibre Channel services class 2 and 3.
- LUN Masking, which is an Initiator-based LUN masking for storage traffic isolation.
- FCP SCSI initiator and target operation.
- Full-duplex operation.
- N_Port Id Virtualization (NPIV) allows multiple host initiator N_Ports to share a single physical N_Port, dramatically reducing SAN hardware requirements.

The IBM Flex System FC5052 2-port 16Gb FC Adapter has the following features:

- 2-port 16 Gb Fibre Channel adapter
- Single-ASIC controller using the Emulex XE201 design
- Up to 1.2 million IOPS on a single port
- Auto-negotiate to 16 Gb, 8 Gb, or 4 Gb
- PCIe Express 2.0 x8 host interface (5 GTps)
- MSI-X support
- IBM Fabric Manager support
- Common driver model with the CN4054 10Gb Ethernet, EN4054 10Gb Ethernet, and FC3052 8Gb FC adapters

Figure 8 shows the IBM Flex System FC5052 2-port 16Gb FC Adapter.

Figure 8 IBM Flex System FC5052 2-port 16Gb FC Adapter

The IBM Flex System FC5054 4-port 16Gb FC Adapter has the following features:

- A 4-port 16 Gb Fibre Channel adapter.
- A dual-ASIC (FC5024) controller that uses the Emulex XE201 design, which allows for logical partitioning on Power Systems compute nodes.
- Up to 1.2 million IOPS on a single port.
- Auto-negotiate to 16 Gb, 8 Gb, or 4 Gb.
- Two PCIe Express 2.0 x8 host interfaces (each 5 GTps), one for each ASIC.
- The ASICs are treated as separate devices by the driver. There are no shared resources (that is, no PCIe bridge) between ASICs.
- Each ASIC has its own firmware resources.
- MSI-X support.
- Common driver model with the CN4054 10Gb Ethernet, EN4054 10Gb Ethernet, and FC3052 8Gb FC adapters.
- IBM Fabric Manager support.

Figure 9 shows the IBM Flex System FC5054 4-port 16Gb FC Adapter.

Figure 9 IBM Flex System FC5054 4-port 16Gb FC Adapter

For more information, see *IBM Flex System FC5052 2-port and FC5054 4-port 16Gb FC Adapters*, TIPS1044, which is available at the following website:

http://www.redbooks.ibm.com/abstracts/tips1044.html

Deployment scenarios of 16 Gb FC host connectivity

Even if your storage system does not have 16 Gb Fibre Channel ports, in certain cases 16 Gb FC host connectivity can provide significant advantages over the 8 Gb FC deployments. These advantages are:

- Lower acquisition costs
- Simplified cabling
- Lower operational costs
- Lower maintenance and support costs

The following subsections provide ideas about where to deploy 16 Gb FC host connectivity using the unique capabilities of the IBM Flex System offerings and the potential benefits of 16 Gb speeds compared to 8 Gb speeds. The following scenarios are described:

- OLAP databases
- Video on Demand
- Mixed virtualized workloads
- Dual VIOS configuration
- Disk and tape traffic separation
- x222 Fibre Channel connectivity

OLAP databases

Data warehouses are commonly used with online analytical processing (OLAP) workloads in decision support systems, for example, financial analysis. Unlike OLTP, where transactions are typically relatively simple and deal with small amounts of data, OLAP queries are much more complex and process large volumes of data. By its nature, the OLAP workload is sequential read-intensive and throughput-intensive.

OLAP databases are normally separated from OLTP databases, and OLAP databases consolidate historical and reference information from multiple sources. Queries are submitted to OLAP databases to analyze consolidated data from different points of view to make better business decisions in a timely manner.

For OLAP workloads, it is critical to have a fast response time to ensure that business decisions support an organization's strategy and are made in a timely manner in response to changing market conditions. Delays might significantly increase business and financial risks; therefore, storage I/O capabilities must match the performance of other server subsystems to ensure that queries are processed as quickly as possible.

For illustration purposes, consider the following scenario. Multiple business analysts need to evaluate current business performance and discover new potential opportunities. They submit OLAP queries, and their queries need to cumulatively process 1 TB of data.

The 16 Gb FC OLAP solution consists of the following components:

- IBM Flex System x440 Compute Node that is used as an OLAP server
- Two IBM Flex System FC5022 or FC5052 2-port 16 Gb FC Adapters that are installed in the x440
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node

These components are installed in IBM Flex System Enterprise Chassis. All connectivity is provided by the chassis midplane and no cables are needed. V7000 Storage Node has eight 8 Gb FC links to the FC5022 switches (four links per switch).

Figure 10 shows the OLAP solution.

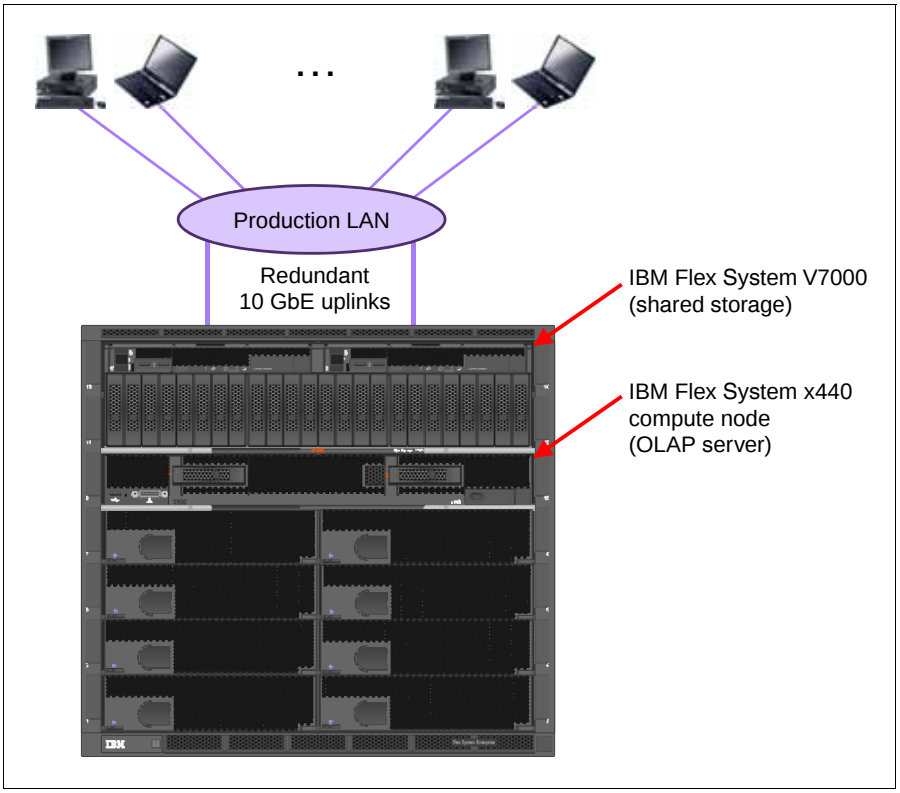

Figure 10 OLAP solution

With 16 Gb FC connectivity, 1 TB of data can be transferred in less than three minutes, assuming that all four 16 Gb ports on the x440 are used and the theoretical bandwidth of 16 Gb FC is 1.6 GBps.

With 8 Gb FC connectivity, 1 TB of data can be transferred in slightly more than five minutes, assuming that all four 8 Gb ports on the x440 are used and the theoretical bandwidth of 8 Gb FC is 0.8 GBps.

In dense deployments, such as IBM Flex System where the number of expansion card slots is limited, higher speed technology provides clear benefits - queries can be completed much faster compared to lower speed technologies. Also, server utilization increases, helping to get faster return on investments. Table 12 summarizes the characteristics of the two approaches.

Table 12 16 Gb versus 8 Gb Fibre Channel: OLAP workload scenario

Characteristic	16 Gb FC	8 Gb FC	16 Gb FC advantage
Number of servers	1	1	
Number of FC adapters	2	2	
Number of storage systems	1	1	
Storage location	Internal	Internal	

Characteristic	16 Gb FC	8 Gb FC	16 Gb FC advantage
Throughput, GBps	1.6	0.8	Faster query execution, more efficient server utilization, more queries per minute
Transfer time, min.	2.6	5.2	

Video on Demand

Although video on demand is traditionally sequential throughput-intensive workload, in a multi-user environment, where every user receives their own data stream watching different content or even the same content with some delay (for example, a recently published new movie), the workload becomes randomized, which requires faster response time to ensure a better user experience and smoother video playback. In general, video streaming applications use 64 KB or more I/O blocks to interact with the storage system.

Video libraries require a significant amount of storage space and sufficient throughput. This can be achieved using a tiered storage design approach with automated tier management capabilities. With such an approach, movie libraries are on nearline (Tier 2) storage, and the movies being watched are placed on the online (Tier 0 or Tier 1) storage.

Consider the following scenario:

A provider of on demand video content has 50,000 subscribers, and 10,000 movies in their video library. 7,500 subscribers are active at the same time, and they watch 1,000 videos simultaneously. The provider uses SD video that requires about 3 Mbps or 0.4 MBps per stream, and the average movie size is 2 GB.

The required nearline storage capacity is 20 TB, and the required online storage capacity is 2 TB.

To meet SLA requirements for 7,500 concurrent users, a fully redundant infrastructure is implemented, including network and storage connections and servers.

Let us assume that a single IBM Flex System x240 server used in the scenario can handle 5,000 simultaneous video streams. The total throughput for 5,000 concurrent streams is 2 GBps; however, the maximum number of actual concurrent connections is limited by the network throughput or storage throughput (whichever is lower).

With 8 Gb FC, the throughput is limited to 0.8 MBps; therefore, the number of concurrent users are limited to 2,000. With 16 Gb FC, the throughput is 1.6 GBps, which corresponds to 4,000 users; however, if only one 10 GbE network link will be used (dual-port adapter is configured in a failover team), the maximum number of concurrent users are 3,000 (because the maximum theoretical throughput of 10 Gb Ethernet is 1.2 GBps). With 4-port 10 GbE adapters, x240 can support up to 4,000 concurrent users.

The solution consists of Video on Demand (VoD) servers that are equipped with dual-port 10 Gb Ethernet adapters and 8 Gb or 16 Gb FC adapters. VoD servers are connected to the shared storage through a SAN fabric. Figure 11 shows the solution components.

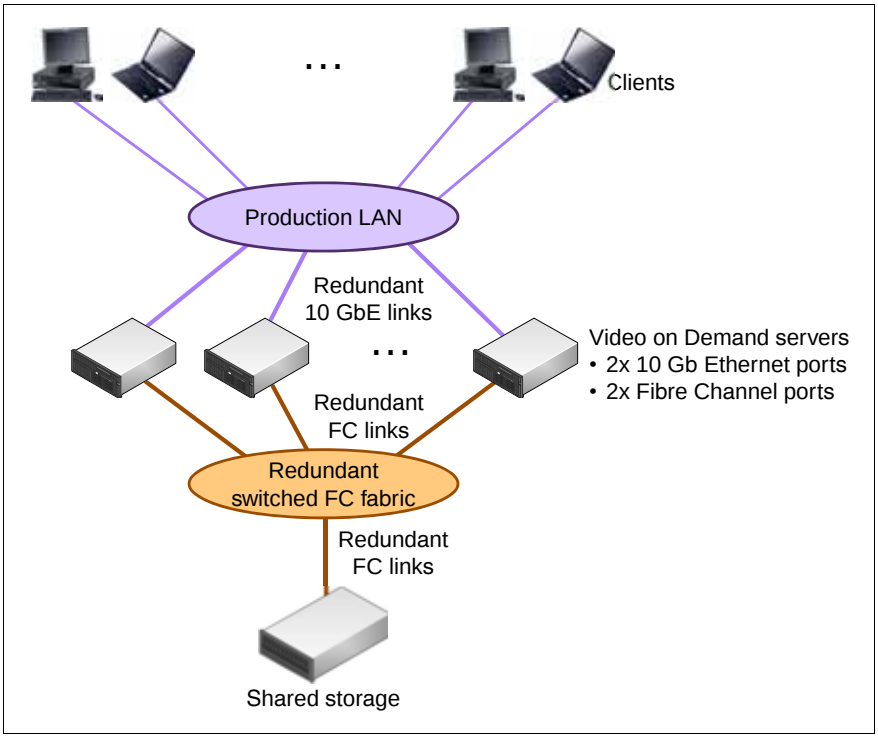

Figure 11 Video on Demand solution components

The IBM Flex System 16 Gb FC solution components that are used in the Video on Demand solution are:

- IBM Flex System x240 Compute Nodes with the following features:
  - Integrated dual-port 10 GbE LOM controllers
  - IBM Flex System FC5022 or FC5052 2-port 16 Gb FC adapters
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node
- IBM Flex System V7000 Expansion Node

The total number of the x240 compute nodes with 16 Gb FC connectivity is four: three nodes are required to support 7,500 concurrent users (up to 3,000 users per one node), and an additional node provides N+1 server redundancy to meet SLA requirements.

With 8 Gb FC adapters, the total number of servers increases to five.

Figure 12 and Figure 13 respectively, show the VoD solutions with 8 Gb and 16 Gb FC connectivity.

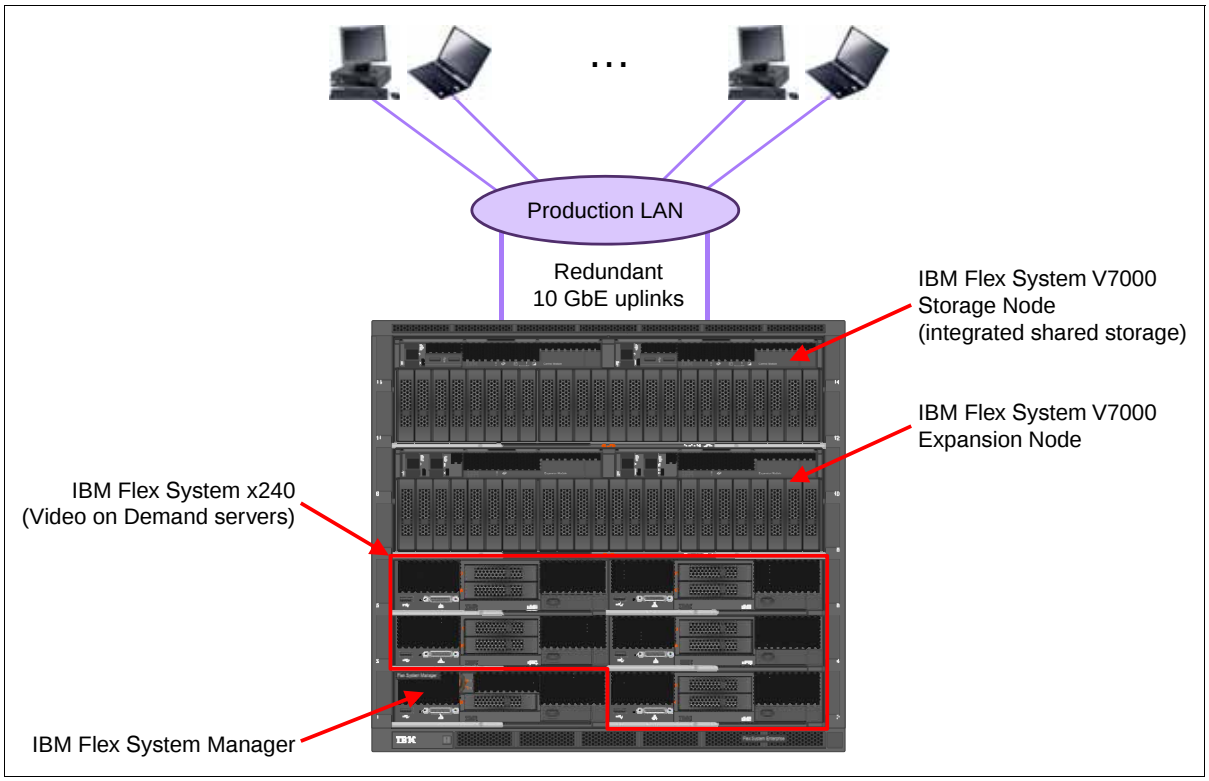

Figure 12 IBM Flex System Video on Demand solution with 8 Gb FC connectivity

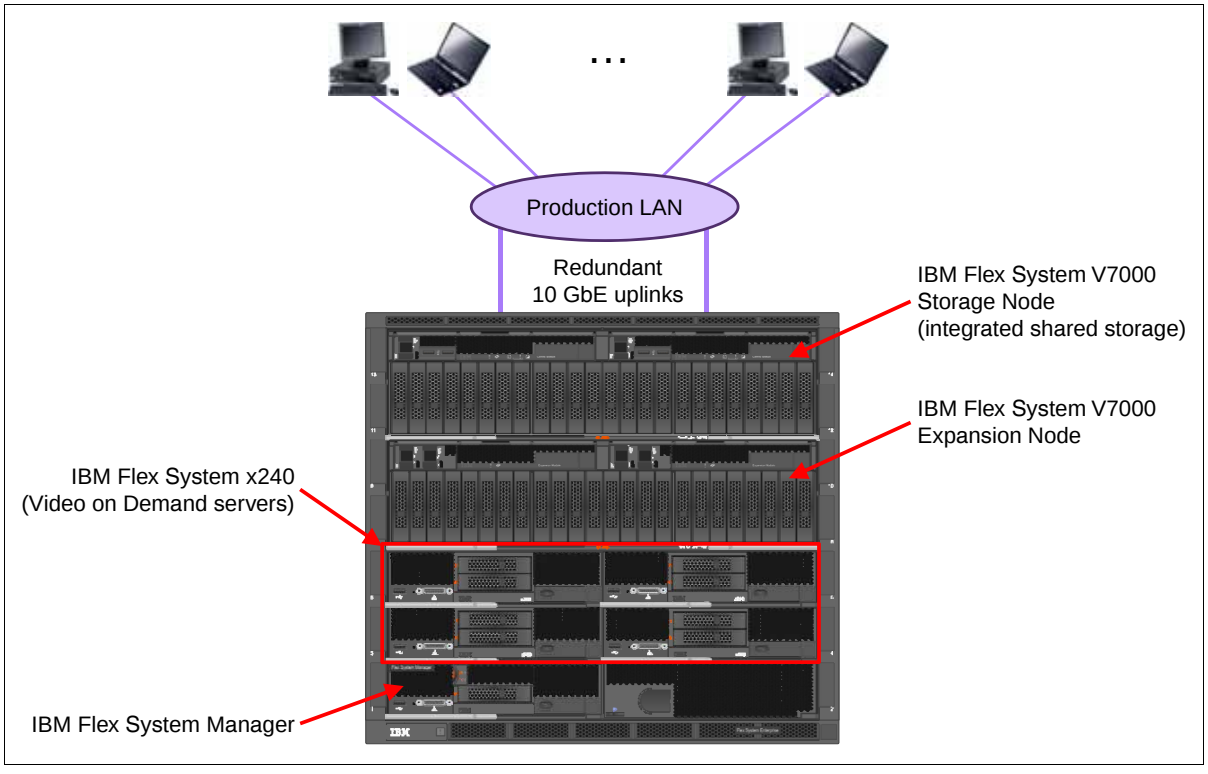

Figure 13 IBM Flex System Video on Demand solution with 16 Gb FC connectivity

Table 13 summarizes the characteristics of these solutions and benefits that are provided by 16 Gb FC.

Table 13 16 Gb versus 8 Gb Fibre Channel: Video on Demand scenario

Characteristic	16 Gb FC	8 Gb FC	16 Gb FC advantage
Number of servers	4	5	Up to 20% fewer components to acquire and maintain, simplified cabling, higher reliability, and easier management
Number of FC adapters	4	5	
Number of storage systems	1	1	
Number of storage expansion enclosures	1	1	
Storage location	Internal	Internal	
Throughput in GBps	3.0	3.0	50% higher number of users per server, and more efficient server utilization
Number of users per server	3000	2000	
Power and cooling costs	Potentially lower	Potentially higher	Lower power and cooling costs
Acquisition costs	Potentially lower	Potentially higher	Lower acquisition costs
Total cost of ownership	Potentially lower	Potentially higher	Lower total cost of ownership

To increase server utilization even more (up to 4,000 users per server), you can install 4-port 10 GbE adapters and further decrease the number of servers from four to three, as shown in Figure 14.

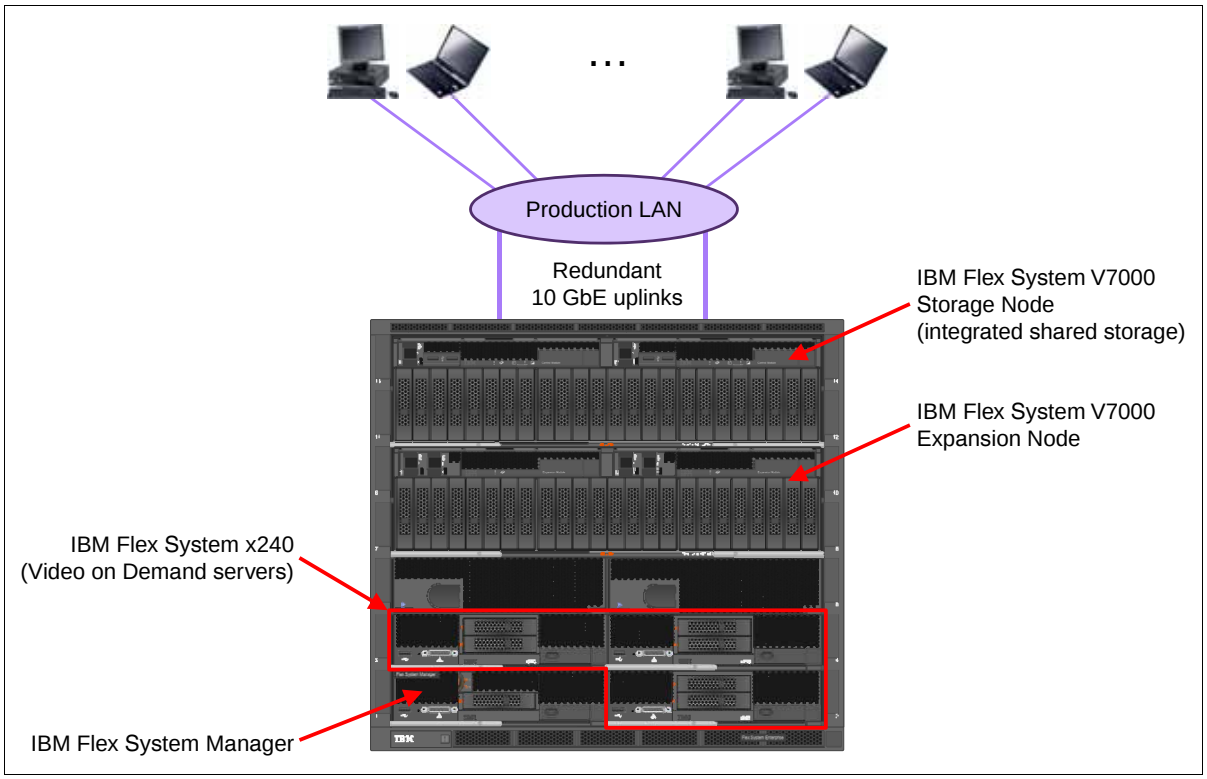

Figure 14 Flex System VoD solution with 16 Gb FC connectivity and four 10 GbE ports per node

Mixed virtualized workloads

In a mixed virtualized environment, we can observe different kinds of workloads and storage access patterns, even on the same physical server.

Consider the following scenario as an example:

There are 44 servers that must be consolidated in a virtualized environment. The server types are the following ones:

- Web server
- Database server
- Streaming media server
- File server
- Collaboration/email server

Table 14 shows the server characteristics and requirements.

Table 14 Virtual machine storage throughput requirements

Server type	Throughput	Quantity	Total throughput
Web server	50 MBps	20	1 GBps
Database server	50 MBps	12	0.6 GBps
Streaming media server	250 MBps	4	1 GBps
File server	100 MBps	4	0.4 GBps
Email server	50 MBps	4	0.2 GBps

The total throughput that is required is 3.2 GBps, and each server hosts a mix of all workloads that are grouped into four sets, and each set consists of five web servers, three database servers, one streaming media server, one file server, and one email server. Assume that IBM Flex System x240 server has sufficient processing capacity (processor, memory, and I/O) to host up to two virtual machine sets.

The solution consists of virtualized servers that are equipped with dual-port 10 Gb Ethernet adapters and 8 Gb or 16 Gb FC adapters. These servers are connected to the shared storage using a SAN fabric.

The solution is shown in Figure 15.

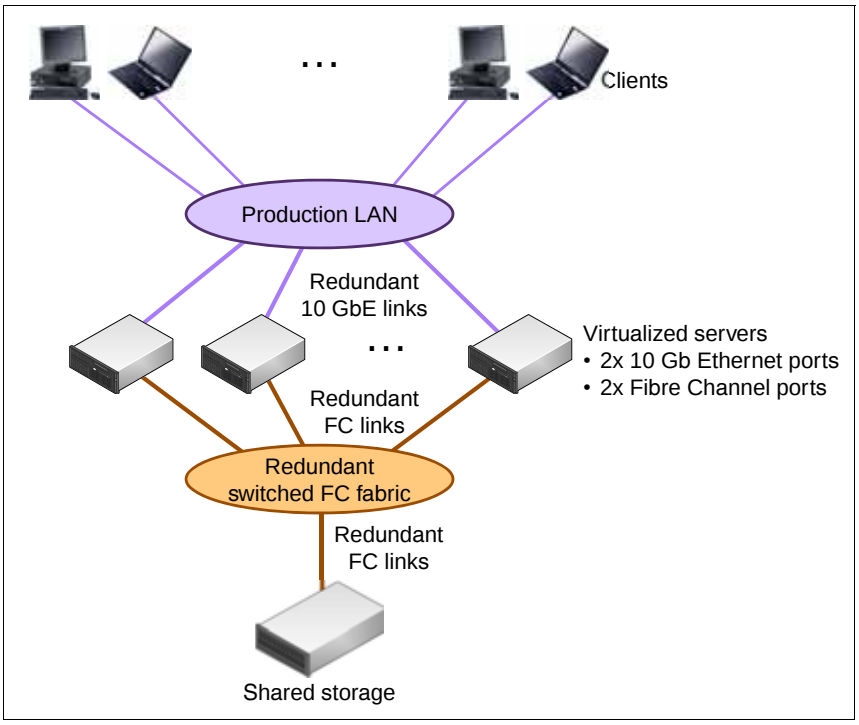

Figure 15 Virtualization solution components

The IBM Flex System 16 Gb FC virtualization solution components are:

- IBM Flex System x240 Compute Nodes with the following features:
  - Integrated dual-port 10 GbE LOM controllers
  - IBM Flex System FC5022 or FC5052 2-port 16 Gb FC Adapters
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node

The total number of x240 compute nodes featuring 16 Gb FC connectivity is three: two nodes are required to support 44 virtual servers (22 per server), and an additional node provides N+1 server redundancy to meet SLA requirements.

With 8 Gb FC adapters, the total number of servers increases to five (11 virtual machines per server plus an additional N+1 redundant node).

Figure 16 and Figure 17, respectively, show the IBM Flex System based virtualization solutions with 8 Gb FC and 16 Gb FC connectivity.

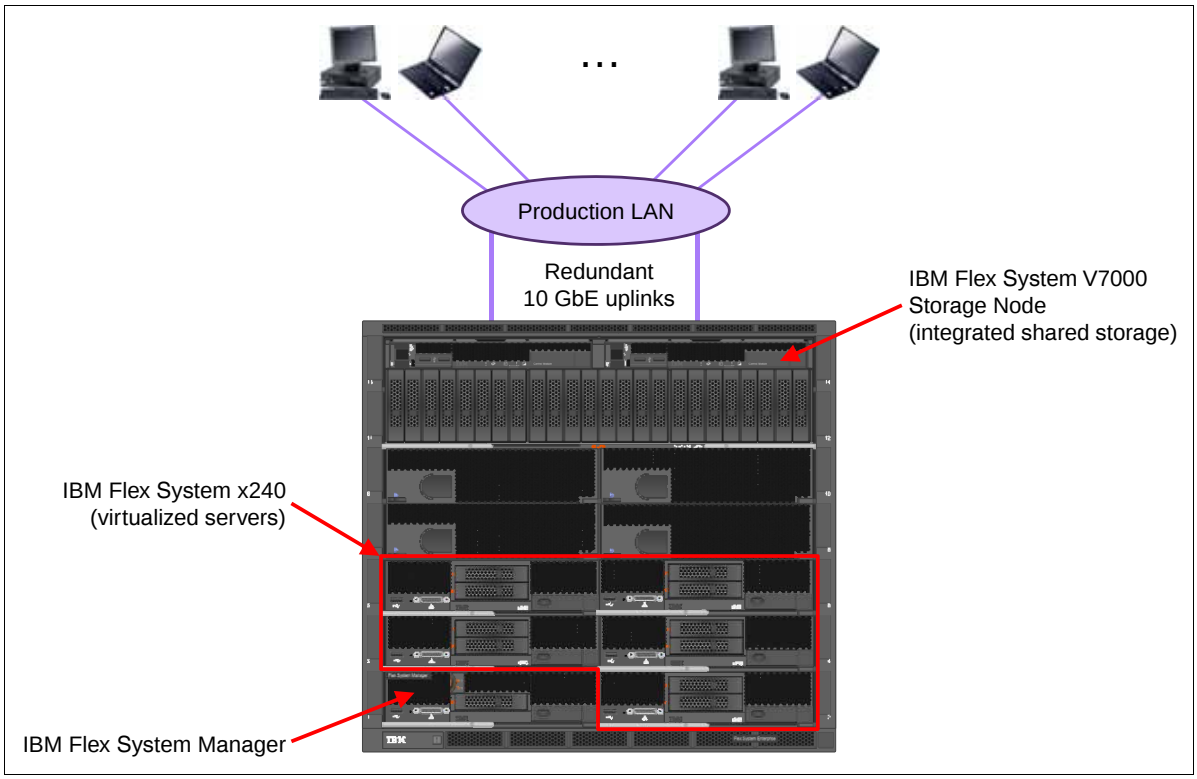

Figure 16 IBM Flex System virtualization solution with 8 Gb FC connectivity

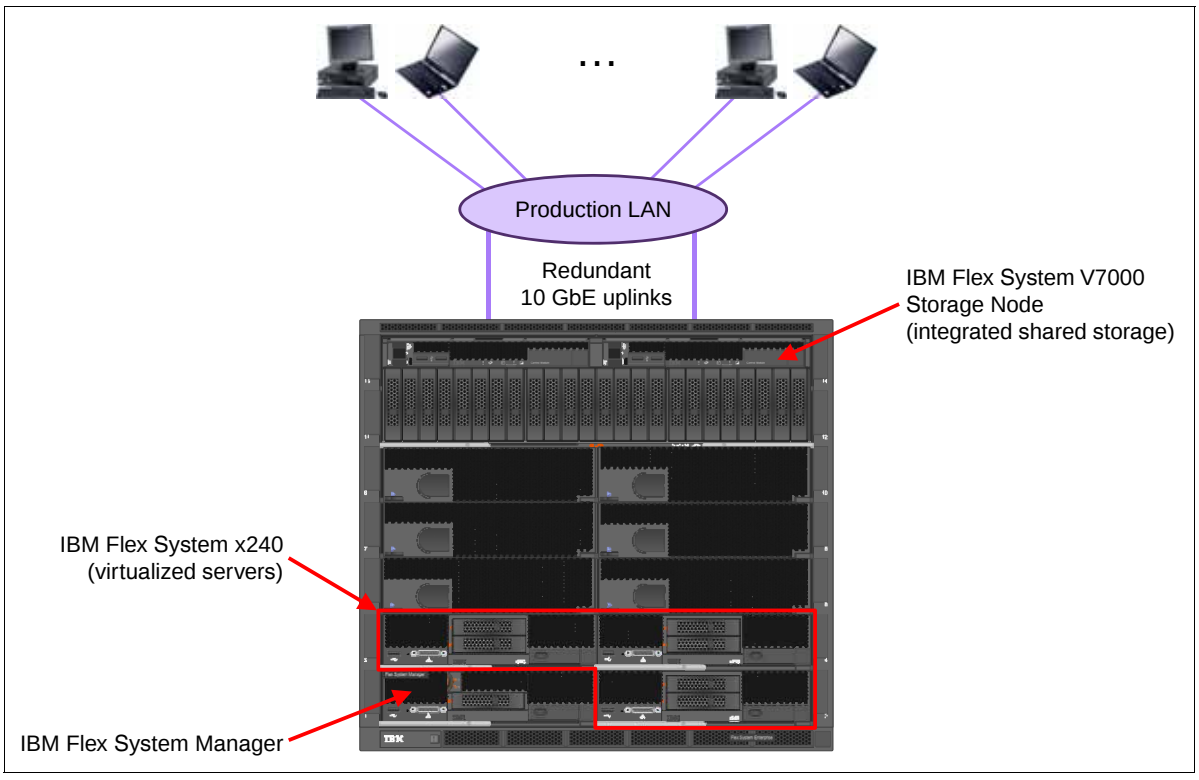

Figure 17 IBM Flex System virtualization solution with 16 Gb FC connectivity

Table 15 summarizes the characteristics of these solutions and benefits that are provided by 16 Gb FC.

Table 15 16 Gb versus 8 Gb Fibre Channel: Mixed virtualized workload scenario

Characteristic	16 Gb FC	8 Gb FC	16 Gb FC advantage
Number of servers	3	5	Up to 40% fewer components to acquire and maintain, simplified cabling, higher reliability, and easier management
Number of FC adapters	3	5	
Number of storage systems	1	1	
Storage location	Internal	Internal	
Throughput in GBps	3.2	3.2	Twice the higher number of VMs per server and more efficient server utilization
Number of VMs per server	22	11	
Power and cooling costs	Potentially lower	Potentially higher	Lower operational costs
Acquisition costs	Potentially lower	Potentially higher	Lower acquisition costs
Total cost of ownership	Potentially lower	Potentially higher	Lower total cost of ownership

Dual VIOS configuration

In scalable, highly available IBM PowerVM® environments, multiple Virtual I/O Servers (VIOSes) provide the performance and redundancy that is expected for those types of environments. Redundancy can be built into the Virtual I/O Server itself by using redundant physical adapters that are combined with MPIO or LVM mirroring for storage devices and Link Aggregation for network devices.

In a dual Virtual I/O Server configuration, a virtual SCSI and Shared Ethernet Adapter can be configured in a redundant fashion allowing system maintenance, such as reboot, software updates, or even reinstallation, to be performed on a Virtual I/O Server without affecting the virtual I/O clients. With correct infrastructure planning and implementation, maintenance can be performed on a Virtual I/O Server *and* any external device that the VIOS connects to, such as a network or SAN switch, which removes the layer of physical resource dependency. This situation results in improved uptime and reduced system administration efforts for the client partitions, and this is the main reason to implement two Virtual I/O Severs.

One of the possible options to configure dual VIOS is to host both VIOSes on the external Fibre Channel storage. In such a case, VIOS is allocated I/O adapters resources at an ASIC level to retain adapter level resiliency, so dual-ASIC adapters are required. An FC5054 4-port 16 Gb FC adapter satisfies this requirement.

The solution topology for dual VIOS storage connectivity is shown in Figure 18.

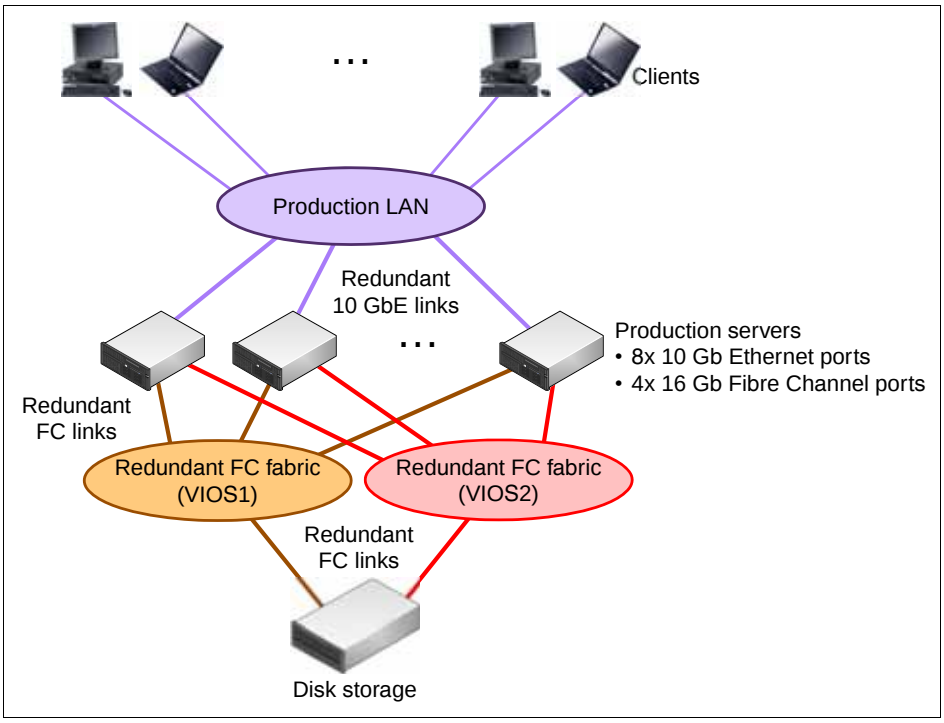

Figure 18 Dual VIOS storage connectivity topology

The IBM Flex System dual VIOS solution components include the following ones:

- IBM Flex System p270 Compute Nodes with the following features:
  - IBM Flex System EN4054 4-port 10 GbE adapters
  - IBM Flex System FC5054 4-port 16 Gb FC Adapters
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node

The Flex System dual-VIOS sample solution is shown in Figure 19.

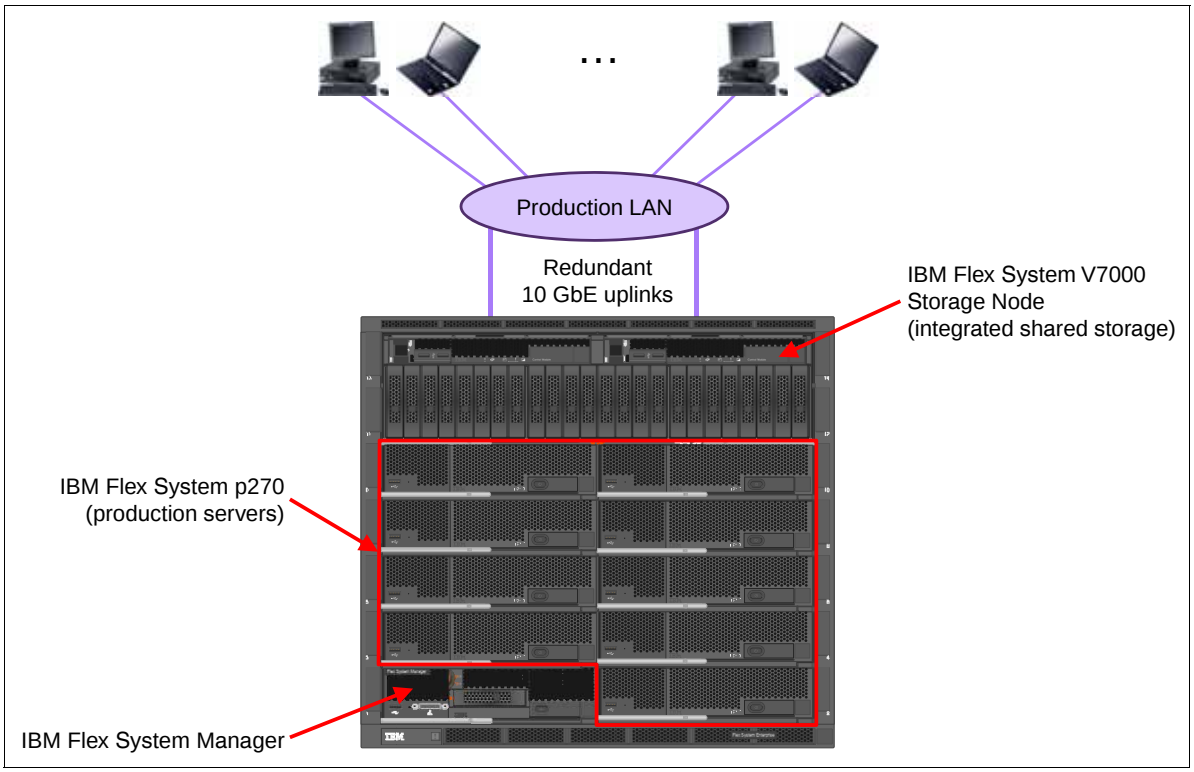

Figure 19 IBM Flex System dual VIOS solution

FC5022 switches that are used in the solution support the Dynamic Ports on Demand (DPOD) feature. With DPOD, ports are licensed as they come online. With the 24-port 16 Gb SAN switch, the first 24 ports reporting (on a first-come, first-served basis) on boot-up are assigned licenses. These 24 ports can be any combination of external or internal Fibre Channel (FC) ports. After all licenses are assigned, you can manually move those licenses from one port to another.

The total number of FC5022 licensed ports that are required in this solution is 44 (22 per FC5022 switch):

- Nine p270 compute nodes with four FC ports per node
- Two 8 Gb FC host interface cards with four ports each in the V7000 Storage Node

As a result, two 24-port FC5022 switches provide a sufficient number of port licenses to implement the solution without a need to buy additional upgrades.

Disk and tape traffic separation

Tape backup solutions are an essential part of a company's overall data backup and recovery strategy. One of the tape backup approaches is a LAN-free backup where the SAN fabric is used for the backup data flow instead of LAN, and LAN is used only for passing control information between the backup server and agents. LAN-free backup can save network bandwidth for network applications, thus providing better network performance. The backup agent transfers backup data from the disk storage directly to the tape storage during LAN-free backup.

The LAN-free backup architecture is shown in Figure 20.

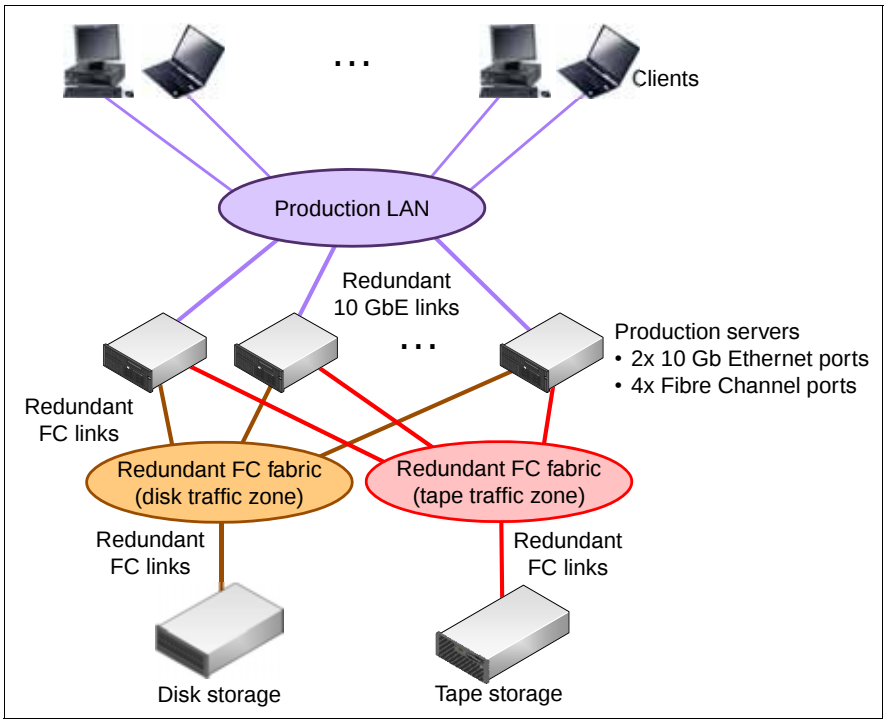

Figure 20 LAN-free backup solution architecture

It is common to separate disk and tape traffic to avoid potential performance issues, but in a fully redundant configuration, such an approach requires four FC ports on a server. In dense server environments such as blade servers, this requirement can represent a challenge because I/O capabilities of the blade servers are often limited.

With IBM Flex System, you can use an FC5054 4-port 16 Gb FC adapter together with FC5022 16 Gb SAN Scalable Switches to build fully redundant storage infrastructure with disk and tape traffic that is isolated from each other.

Here are the IBM Flex System 16 Gb LAN-free tape backup solution components:

- IBM Flex System x240 Compute Nodes with the following features:
  - Integrated dual-port 10 GbE LOM controllers
  - IBM Flex System FC5054 4-port 16 Gb FC Adapters
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node
- IBM System Storage® tape (TS3200 is used in the example that is shown in Figure 21 on page 38).

This solution is shown in Figure 21.

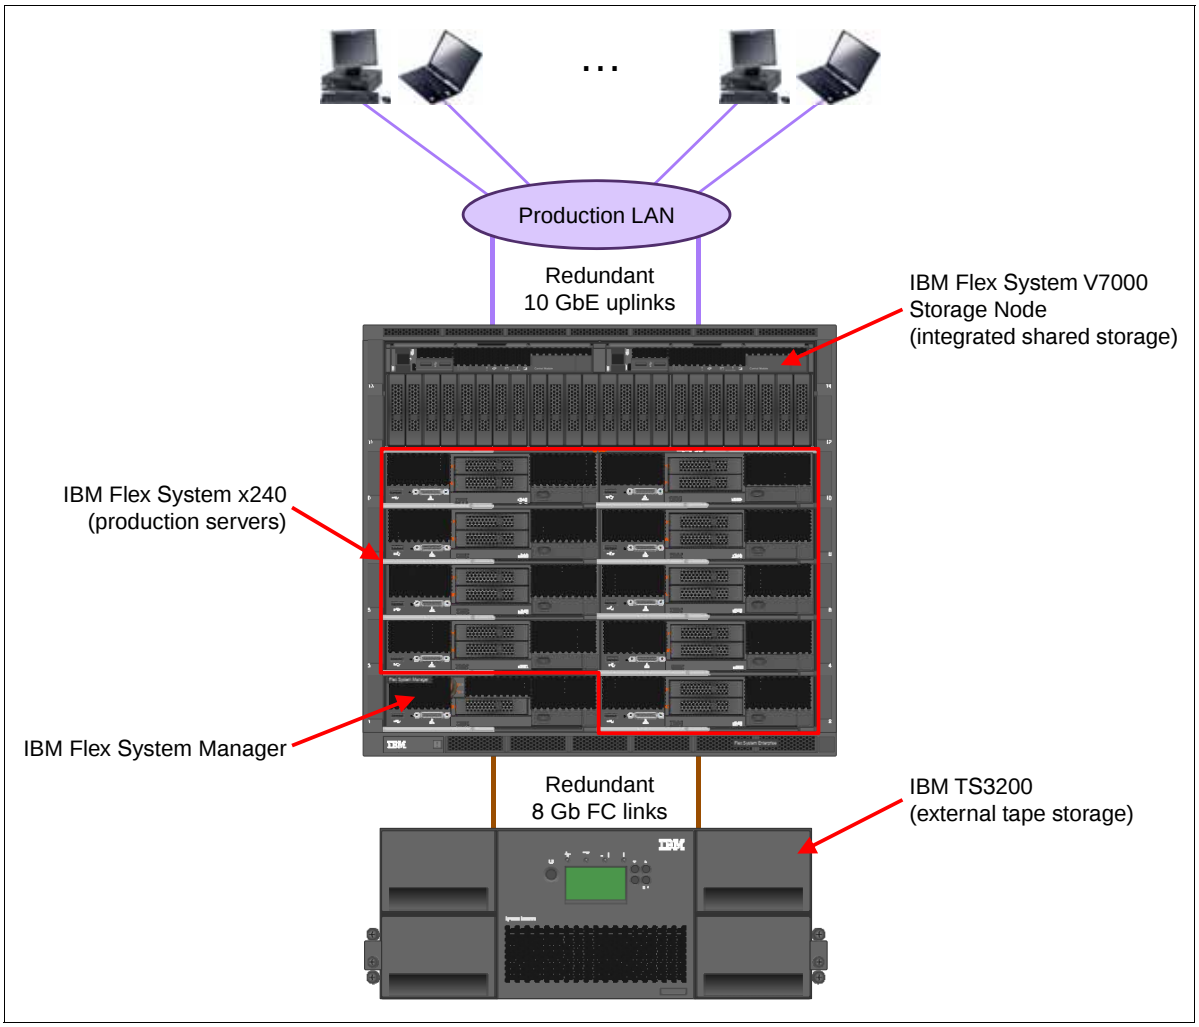

Figure 21 Flex System LAN-free backup solution

In the scenario that is shown in Figure 21, nine x240 compute nodes are equipped with the FC5054 4-port adapters, a V7000 Storage Node has two four-port FC host interface cards, and two external ports on each FC5022 switch are connected to the external TS3200 tape library. Therefore, the total number of FC ports that are required is 48 (24 ports per FC5022 switch).

The DPOD feature of the FC5022 switches allows FC5054 adapters, an integrated V7000 storage node, and an external tape library to be connected to the 24-port FC5022 switches without a need to buy additional port upgrades, which helps provide cost-effective storage connectivity.

x222 Fibre Channel connectivity

The IBM Flex System x222 Compute Node is a high-density offering that maximizes the computing power that is available in the data center. With a balance between cost and system features, the x222 is an ideal platform for dense workloads, such as virtualization.

The x222 occupies one half-wide compute node bay in the Flex System Enterprise Chassis, and it consists of two independent compute servers (upper server and lower server), which requires that a unique I/O architecture be implemented. The x222 supports one I/O adapter that is shared between the two servers and is routed to the I/O modules that are installed in bays 3 and 4 of the chassis.

When Fibre Channel connectivity is required for the x222, the FC5024D adapter is used. The FC5024D is a four-port adapter where two ports are routed to each server. Port 1 of each server is connected to the switch in bay 3, and Port 2 of each server is connected to the switch in bay 4. To make full use of all four ports, you must install FC5022 16 Gb FC switches with a sufficient number of licensed ports in both I/O module bays.

The x222 16 Gb FC I/O connectivity topology with FC5022 switches is shown in Figure 22.

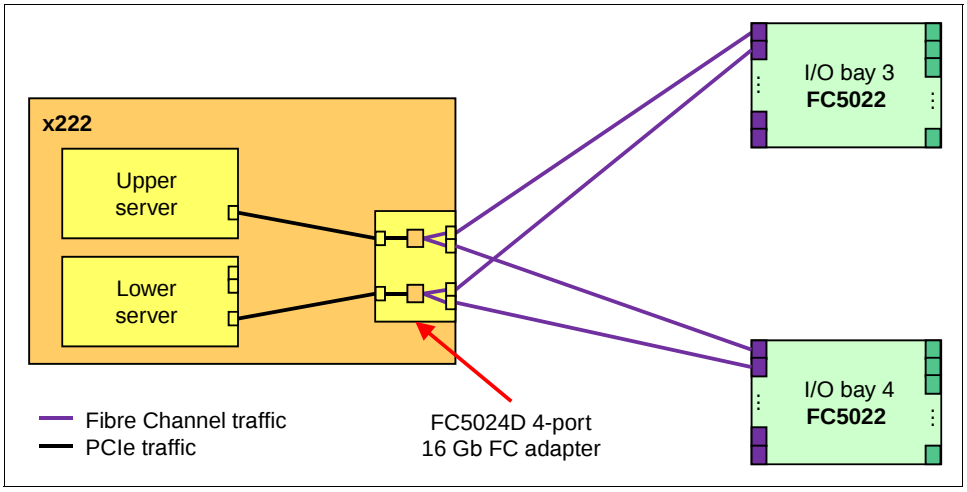

Figure 22 The x222 16 Gb FC I/O connectivity topology

A typical IBM Flex System x222 16 Gb FC storage connectivity solution consists of the following components:

- IBM Flex System x222 Compute Nodes with the following features:
  - Integrated dual-port 10 GbE LOM controllers
  - IBM Flex System FC5024D 4-port 16 Gb FC Adapters
- Two IBM Flex System FC5022 24-port 16 Gb SAN Scalable Switches
- IBM Flex System V7000 Storage Node

The solution is shown in Figure 23.

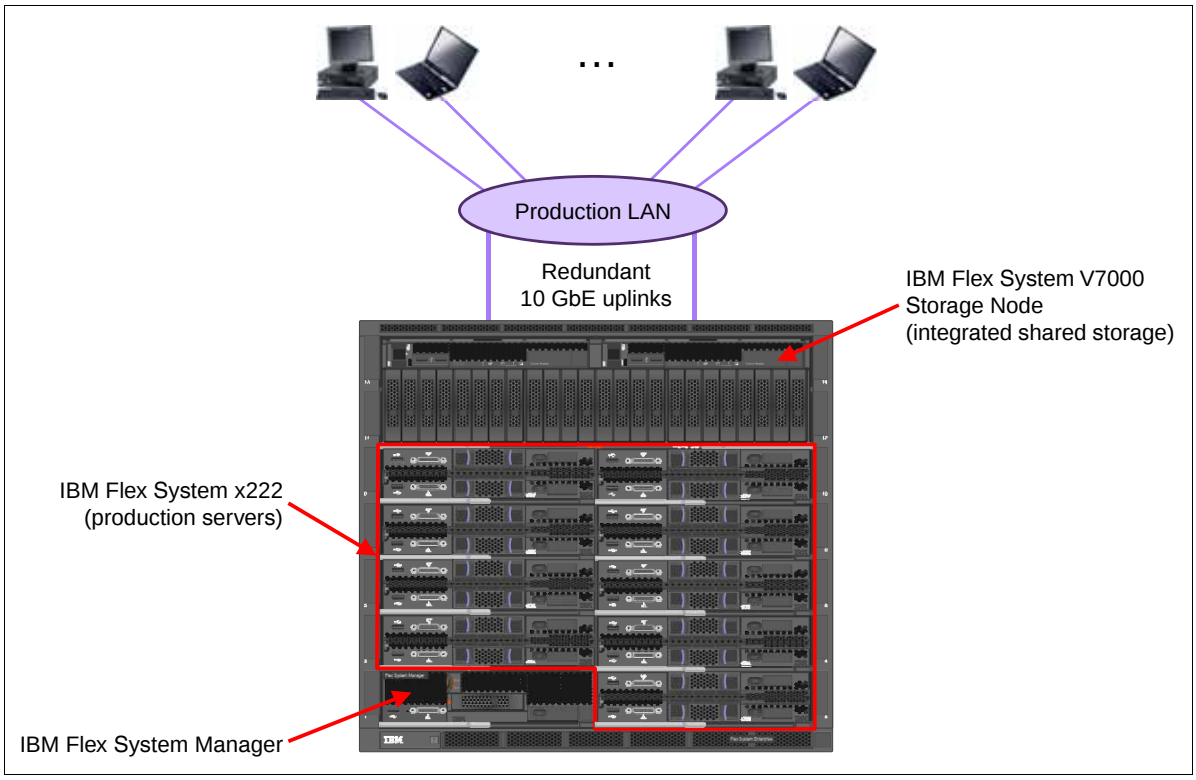

Figure 23 IBM Flex System x222 16 Gb FC storage connectivity

The total number of FC5022 licensed ports that are required in this solution is 44 (22 per FC5022 switch):

- Nine x222 compute nodes with four FC ports per node bay
- Two host interface cards (HICs) on the V7000 with four 8 Gb FC ports on each HIC

Therefore, 24-port FC5022 16 Gb SAN switches with Dynamic Port on Demand feature satisfy these requirements without a need to buy additional port upgrades, thus providing cost-efficient 16 Gb FC connectivity.

Proof-of-concept scenarios

To analyze performance improvements with 16 Gb Fibre Channel compared to 8 Gb FC, we evaluated two proof-of-concept (PoC) scenarios and evaluated the results that were obtained by running storage I/O benchmarks.

The following scenarios were evaluated:

- Inter-switch links
- Server virtualization

Inter-switch links

In our first scenario, we evaluated how 16 Gb Fibre Channel can help to improve the performance of inter-switch links (ISLs).

We built our configuration using the following components:

- IBM Flex System Enterprise Chassis
- Two EN5022 16 Gb SAN Scalable Switches
- Two x240 compute nodes with FC5022 2-port 16 Gb FC Adapters
- External Brocade 16 Gb FC Top-of-Rack (ToR) switch
- IBM Storwize® V7000 external storage

The solution topology is shown in Figure 24.

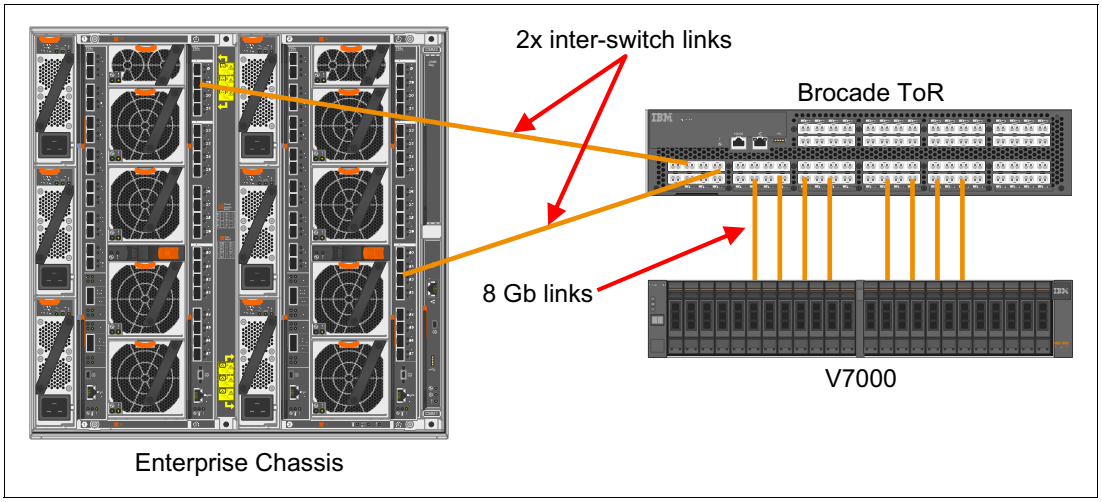

Figure 24 ISL PoC topology

Each of the FC5022 SAN switches in the Enterprise Chassis is connected to the ToR switch using a single ISL link. The V7000 is connected to the external ToR switch using eight 8 Gb FC links. The KGen tool is used to generate a sufficient storage I/O workload from the x240 compute nodes. The results are shown in Table 16.

Table 16 Storage performance measurements: ISLs

Measurement	8 Gb FC	16 Gb FC	Improvement ratio
Throughput	0.76 GBps	1.4 GBps	1.84x

When the ISL link was operating at 8 Gb, the throughput measured was 0.76 GBps. At 16 Gb, the throughput was 1.4 GBps. A 16 Gb FC helped to increase the performance of a single ISL by more than 80%.

The result that was observed proves that 16 Gb FC can help to reduce the number of ISLs by almost two times.

Server virtualization

In our second scenario, we evaluated how 16 Gb Fibre Channel can help to improve storage I/O performance of the compute node in a virtualized environment.

We built our configuration using the following components:

- IBM Flex System Enterprise Chassis
- Two EN5022 16 Gb SAN Scalable Switches
- Two x240 compute nodes with FC5022 2-port 16 Gb FC Adapters
- IBM Storwize V7000 external storage

The solution topology is shown in Figure 25.

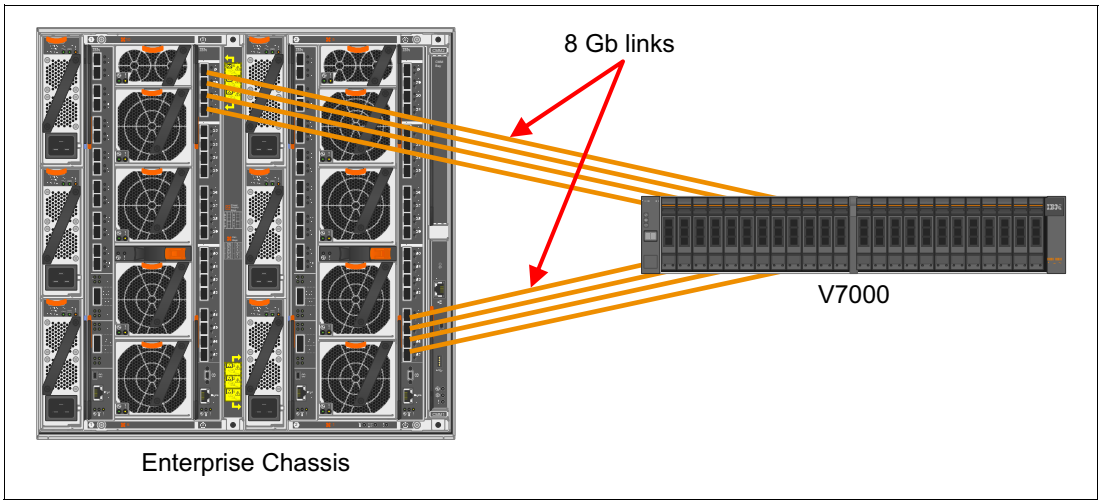

Figure 25 Server virtualization PoC topology

Each of FC5022 SAN switches in the Enterprise Chassis is connected to the V7000 using four 8 Gb FC links. VMware ESXi was deployed on the compute nodes, and KGen was running in a virtual machine (two virtual machines per one compute node). ESXi was configured to balance the storage I/O workload across all available paths. The results are shown in Table 17.

Table 17 Storage performance measurements: Virtualization

Measurement	8 Gb FC	16 Gb FC	Improvement ratio
Throughput	0.7 GBps	1.35 GBps	1.93x

With the host adapters running at 8 Gb speeds, the average host throughput was about 0.7 GBps. At 16 Gb speeds, the average host throughput achieved 1.35 GBps, which is more than 90% improvement compared to 8 Gb.

The result proves that 16 Gb FC can help double storage I/O performance capacity and therefore reduce the number of servers that are required to support specific workloads.

Conclusion

The growth in CPU processing power far exceeds the growth in storage I/O. For this reason, storage I/O is the culprit in major bottlenecks in many high-performance applications. The 16 Gb FC host connectivity can provide up to 1.6 GBps per port for throughput-intensive workloads, making it possible for you to potentially eliminate I/O bottlenecks in a system.

IBM 16 Gb FC solutions support moderate to heavy storage workloads, while providing better performance, availability, scalability, and security compared to converged networks.

With 16 Gb FC, you can potentially deploy fewer number of servers and provide better server utilization, including twice higher density of virtual machines and a higher number of concurrent users, compared to 8 Gb FC solutions.

Integrated 16 Gb FC technology can help to achieve the following benefits:

- Up to 50 - 100% better virtual machine density and higher number of concurrent users due to increased storage bandwidth.
- Up to 20 - 40% fewer servers are required to support the workload that is specified.
- Twice lower number and higher speed of inter-switch links required in scalable SANs
- Higher reliability and availability of the services due to fewer number of components that are used to build the solution.
- Twice faster access to the business critical data.
- Lower acquisition costs due to fewer number of systems and components.
- Shorten ROI time frame and decrease overall TCO with the efficient utilization of server resources and lower power, cooling, and management costs.

Author

This paper was produced by a technical specialist working at the International Technical Support Organization, Raleigh Center.

Ilya Krutov is a Project Leader at the ITSO Center in Raleigh and has been with IBM since 1998. Before he joined the ITSO, Ilya served in IBM as a Run Rate Team Leader, Portfolio Manager, Brand Manager, Technical Sales Specialist, and Certified Instructor. Ilya has expertise in IBM System x®, BladeCenter®, and Flex System products, server operating systems, and networking solutions. He has authored over 150 books, papers, Product Guides, and Solution Guides. He has a Bachelor's degree in Computer Engineering from the Moscow Engineering and Physics Institute.

Special thanks to Khalid Ansari, a member of the IBM Advanced Technical Skills team, who hosted the proof-of-concept project environment, ran the tests, and obtained the benchmarking results for this paper.

Thanks to the following people for their contributions to this project:

Kevin Barnes, Tamikia Barrow, Mary Comianos, Shari Deiana, Cheryl Gera, KaTrina Love, Linda Robinson, Wade Wallace, David Watts
International Technical Support Organization, Raleigh Center

Shekhar Mishra and Mike Talplacido
IBM

Jay Flora, Tony McKenzie, Marcus Thordal, Van Tran, Bill Stelling, Tim Werts, Matt Wineberg
Brocade

Jeff Lin, Tom Boucher, Chris Sorensen
Emulex

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

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 grant 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

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 measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 REDP-4921-01 was created or updated on December 13, 2013.

Send us your comments in one of the following ways:
- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@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.

Redpaper™

Trademarks