PureSystems

IBM Flex System

Going beyond blades

IBM

Highlights

Networking capabilities designed for virtualization and cloud

- High-performance Ethernet and converged options are available. Clients also have capability to implement end-to-end 1 Gb, 10 Gb, 40 Gb Ethernet and Fibre Channel over Ethernet (FCoE) switches to fit into your current infrastructure while supporting next-generation capabilities
- Fibre Channel switches ranging from low-cost 8 Gb solutions to highly scalable 16 Gb switches for maximum performance
- Simple management with an integrated fabric managed by a single I/O device, allowing the ability to grow the IBM® Flex System® points of delivery (PODs) supporting up to 252[1] compute nodes

Compute nodes that go beyond blades

- Price-optimized entry nodes to performance-optimized 4- and 8-socket nodes for your most demanding applications
- Available with the highest capability processors and industry-leading memory capability for x86-based compute nodes and the latest POWER7+™ processors for maximum performance and efficiency
- Double-dense design can support the equivalent of 28 nodes per IBM Flex System Enterprise Chassis
- Storage capabilities optimized for efficiency and performance
- Optional IBM Flex System V7000 Storage Node offers storage virtualization, Real-time Compression™ to reduce capacity needs, and IBM Easy Tier® to optimize data placement
- High IOP scalable storage using eXFlash DIMMS for fast, consistent local storage
- Up to 2x more dedicated flash storage for compute nodes compared to Dell[2]

Management integration with IBM Flex System Manager™

- Single point of control: management of physical and virtual compute, storage and networking resources from a single management console, including x86 and POWER environments
- Smart provisioning: simple and repeatable hardware configuration and bare metal provisioning for server nodes, and five click creation and zoning of storage volumes.
- Upward integration: integrate hardware management functions with VMware vCenter or Microsoft System Center.
- Mobile management: monitor your IBM PureFlex® System and IBM Flex System infrastructure from anywhere using iOS, Android and Blackberry mobile devices.
- Simple administration: Flex System Manager fuel gauge calculates remaining systems management capacity and makes recommendations for FSM performance optimization.
- Increased scalability: manage up to 16 Flex System virtualized chassis, 224 compute nodes, and a total of 5000 end points from a single Flex System Manager.

You know your business and your IT needs. You want to build a solution tuned for your applications and services. You need the most advanced blade technology available, and a flexible architecture to support your needs today—and tomorrow. And you need simple, integrated management to keep operational costs down. But you also want a no-compromise design and long-term investment protection.

IBM Flex System compute, storage, management, and networking offer performance and capacity to support your most demanding applications. IBM Flex System technology delivers the performance and flexibility you need now and is designed to support multiple generations of future technologies.

Chassis

The IBM Flex System Enterprise Chassis is a simple, integrated infrastructure platform for your system that supports a mix of compute, storage, systems management, and networking resources to meet the demands of your applications. The solution is easily scalable by adding chassis with the required nodes. With Flex System Manager, multiple chassis can be monitored from a single point of control. The 14-node, 10U chassis delivers high-speed performance complete with integrated servers, storage and networking. This flexible chassis is designed to deploy simply now and to scale to meet your needs in the future. Clients may also choose between 2500 W and 2100 W power supply modules to optimize their power consumption.

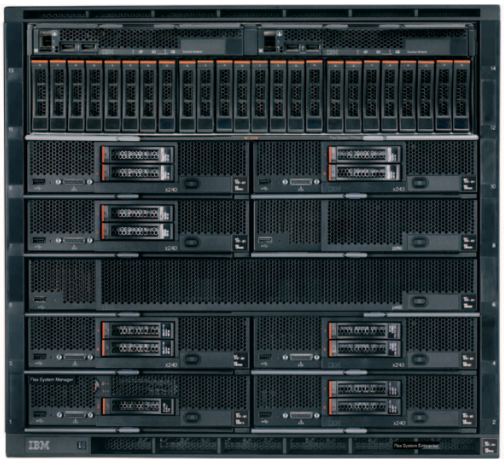

IBM Flex System Enterprise Chassis

Networking capabilities

IBM Flex System Fabric is the interconnect technology for the IBM Flex System. IBM Flex System Fabric optimizes your networking resources for maximum performance, automates network deployment and offers integrated management. The network resources in IBM Flex System are tightly integrated into the system to support virtualization and simple, integrated management. That means you can move from managing a physical network to managing a logical network in a virtualized environment—supporting business services instead of network components. Clients not only have choice of 1 Gb, 10 Gb or 40 Gb network, they can also choose between a full Layer 2/3 switch for highest flexibility to a simple connectivity module for easy setup and management. With integrated management tools based on open standards, these resources are easy to provision and deploy, thereby reducing the cost of managing the virtual fabric.

The new IBM Flex System Interconnect Fabric provides a data center architecture solution designed to offer greater efficiency and automation in network environments. The Flex System Interconnect Fabric offers a solid foundation of compute, network, storage, and software resources in a Flex System POD.

The solution components include; the IBM RackSwitch G8264CS and IBM Flex System SI4093 System Interconnect Modules for data center Ethernet and Fibre Channel network access and fabric interconnect. As a tightly integrated solution, the Flex System Interconnect Fabric presents only one manageable element. All provisioning is accomplished through the primary RackSwitch G8264CS. Once a change is implemented, the primary switch manages implementation of provisioning updates to all devices across the entire Flex System Interconnect Fabric. No manual intervention is required. In a system that can feature up to 20 I/O devices, this results in a 95 percent[3] reduction in the number of elements that must be managed.

In addition to IBM Flex System Fabric, the IBM Flex System networking capabilities provide a choice of multiple fabrics on standards-based protocols. Clients can choose Ethernet (1 Gb, 10 Gb or 40 Gb), Fibre Channel (8 Gb or 16 Gb), FCoE using a pass-thru 10 Gb transit switch or 10 Gb converged switch as well as InfiniBand (QDR or FDR). All modules are built using industry standards and enterprise-class features that have a record of easy interoperation with existing infrastructures. IBM provides extreme flexibility since clients have pay-as-you-grow scalability that allows for the purchase of licenses to activate ports without the complexity of buying, deploying or installing new hardware. Now clients can also leverage flexible port mapping which allows the flexibility to assign internal/external ports as needed.

Ethernet portfolio

For clients looking for 1 Gb, 10 Gb or 40 Gb Ethernet connectivity, the IBM Flex System networking portfolio also offers scalable 1/10 Gb switches and 10 Gb pass-thru modules. For high-end virtualization and high-performance requirements, IBM also offers an end-to-end 40 Gb solution.

Fibre Channel portfolio

Today, Fibre Channel is the dominant choice for storage connectivity. The IBM Flex System portfolio offers both 8 Gb and 16 Gb SAN-connectivity offerings for easy integration with your SAN environment. Advanced licensing features enable advanced SAN functions and monitoring. High-performance, scalable offerings are designed for simple and cost-effective scalability for future growth. And advanced virtualization features enable both physical and virtual SAN setup and management.

The IBM Flex System portfolio offers several 2-port and 4-port Fibre Channel adapters to match the switch offerings. These adapters offer choice to clients in order to have easy interoperability with their existing SAN.

Flex System Fabric EN4093R 10 Gb Scalable Switch

Flex System Fabric Converged offerings

Many clients today are connecting Ethernet and Fibre Channel from their servers upstream into their LAN and SAN. The IBM Flex System Fabric portfolio offers clients several options so they can pick the model that best fits their existing network infrastructure. Clients looking to consolidate multiple chassis can use the EN4093 or SI4093 module in the chassis connected to the IBM G8264CS converged top-of-rack switch. This converged switch can split the FCoE packets to the existing LAN and SAN network. Clients looking to connect their existing LAN and SAN directly to the chassis switch can use the integrated converged switch—CN4093. The IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch allows clients to connect the Flex System chassis to their existing infrastructure without requiring a converged top-of-rack switch. This switch supports multiple protocols such as Ethernet, FCoE, Fibre Channel and iSCSI. With multi-protocol support, it can connect directly to the IBM integrated storage node—offering an integrated solution that is easy to setup and manage—or to an external SAN. When combined with the on-board LOM or the CN4054 Virtual Fabric adapter, this switch offers a cost- effective simple solution that combines the benefits of I/O convergence (FCoE or iSCSI) and Virtual Fabric, while providing 40 Gb uplinks for maximum performance and low latency and support for enhanced features like IBM Virtual Fabric and VMready®.

The IBM Flex System portfolio offers several 2-port and 4-port Ethernet and converged adapters to match the switch offerings. Some of these adapters support Virtual Fabric capability that helps clients reduce cost and complexity.

Infiniband portfolio

High performance computing or financial services solutions require maximum bandwidth and low latency. To support these demands, IBM offers the next-generation InfiniBand solution at both Fourteen Data Rate (FDR) and Quad Data Rate (QDR) speed. The options support Host Channel Adapters (HCA) and switches to connect servers with remote storage and networking devices, and with other servers.

Compute nodes

Taking advantage of the full capabilities of the latest Intel, POWER and POWER7+ processors, IBM Flex System compute nodes are designed to offer the performance you need for your critical applications. With support for a range of hypervisors, operating systems and virtualization environments, the compute nodes provide the foundation for:

- Virtualization solutions
- Virtual desktop solutions
- Business and database applications
- Infrastructure support

The **IBM Flex System x220 compute node** has a versatile, cost-optimized design for infrastructure and entry virtualization workloads. The Flex System x220 compute node features a no-compromise design for processor, memory, and networking resources to allow your business to do more. Features such as automated power management with onboard sensors give you more control over power and thermal levels across the system. These capabilities, combined with memory capacity up to 384 GB are designed to help you get the most out of your systems. Optional Virtual Fabric allows you to take advantage of up to 32 ports of virtual networking capabilities with multiple protocols, including Ethernet, FCoE and iSCSI. The system also allows you to enable features on demand for a more flexible I/O solution.

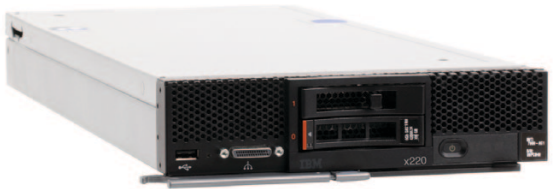

Flex System x220 Compute Node

The new **IBM Flex System x222 compute node** provides a versatile, dense design for VDI, virtualization and cloud workloads. The Flex System x222 compute node delivers this dense design without compromising processor, memory, and networking resources to allow your business to do more.

Each Flex System x222 compute node features two independent twin compute nodes, enabling the equivalent of up to 28 independent compute nodes in a 10U Flex System Enterprise Chassis. There is no need for additional switching hardware or connectivity options, which helps reduce the cost of the overall solution. Finally, clients can still benefit from using Flex System x222 along with other Flex System compute nodes to create a flexible and optimized chassis solution for their diverse workloads.

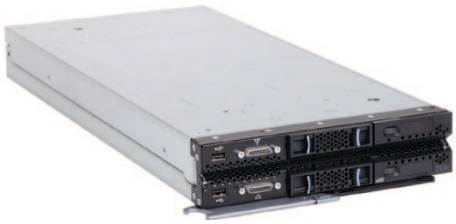

Flex System x222 Compute Node

The **IBM Flex System x240 compute node** is optimized for virtualization, performance and highly scalable I/O designed to run a wide variety of workloads. The Flex System x240 compute node delivers maximum performance—up to 80 percent performance boost over previous generation servers.[4] This helps you get more out of your compute environment for a broad set of workloads. Features such as automated power management with onboard sensors give you more control over power and thermal levels across the system. Offering up to 50 percent more performance per watt than previous generations, the new Flex System x240 is designed to improve energy efficiency.[5] These capabilities, combined with memory capacity up to 768 GB, helps you get the most out of your compute environment.

The **IBM Flex System PCIe Expansion Node** is a component of the IBM Flex System designed to support additional PCI express adapters, fabric mezzanine adapters and next-generation graphics processing units (GPU) to IBM Flex System x220 and x240 compute nodes and expands the compute nodes' capabilities.

Flex System x240 Compute Node

The new **IBM Flex System X6 compute node portfolio** incorporates the sixth generation of IBM enterprise X-Architecture® to help deliver better, more efficient business results using the IBM Flex System. X6 platforms are fast, agile and resilient so they can produce 100 percent faster compute performance than previous-generation systems.[6] The X6 portfolio increases virtualization density, and decreases infrastructure costs and complexity. This enables you to design faster analytics engines, rein in IT sprawl and deliver information with high reliability.

With IBM eXFlash memory-channel storage, the Flex System X6 compute nodes can deliver up to 12.8 TB of ultra-low latency flash storage—unmatched in an x86 blade server.[7] Equipped with the Intel Xeon E7 v2 processors, the Flex System X6 compute nodes can deliver up to 12 TB of memory and 120 cores of processing. The Flex System X6 compute nodes can host business-critical applications, implement large virtual machines or run in-memory databases without compromises in performance, capacity or scalability.

The **IBM Flex System X6 compute node family** consists of three new servers for the Flex System portfolio. The Flex System x280 X6 is a two-socket server optimized for workloads that require a large memory footprint and virtualization performance. The Flex System x480 X6 is scalable to a four-socket system optimized for applications that require more processor cores and memory for larger virtual machines and databases. The Flex System x880 X6 is scalable to an eight-socket system ideal for workloads that demand the highest levels of performance such as analytics and large databases.

IBM Flex System x480 X6

The **IBM Flex System x440 compute node** is ideal for higher-end virtualization, mainstream database deployments, and memory-intensive high performance environments. The Flex System x440 is price-performance optimized and designed to help you match system capabilities and cost to workloads without compromise. This four-socket compute node features no compromise compute, memory and I/O capacity to meet your needs. The Flex System x440 is designed to deliver an outstanding combination of performance, built-in capabilities, and cost-effectiveness, featuring automated power management with onboard sensors to give you more control over power and thermal levels across the system. Combined with memory capacity up to 1.5 TB, the Flex System x440 compute node is designed to help you get the most out of your systems.

Flex System x440 Compute Node

The **IBM Flex System p260 compute node** offers large memory capacity, outstanding performance of the POWER7+ processor, industrial-strength virtualization and workload-optimizing capabilities for small-to-midsize database servers and consolidation of virtualized application servers.

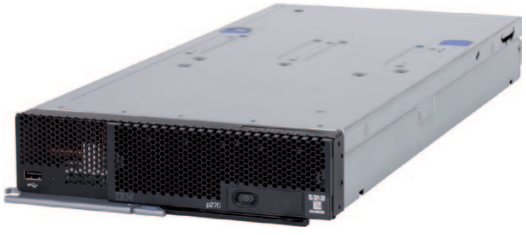

Flex System p270 Compute Node

The new **IBM Flex System p270 compute node** is a POWER7+-based server optimized for virtualization, performance and extraordinary efficiency. The node supports IBM AIX, IBM i or Linux operating environments and is designed to run a wide variety of workloads in your IBM PureFlex System or IBM Flex System solution.

The **IBM Flex System p460 compute node** is an outstanding offering for mid-market clients desiring a high- performance, reliable, secure system that is cloud-enabled and has room for handling their business growth. With excellent virtualization capabilities and the flexibility to run proven solutions from thousands of ISVs that support the AIX®, IBM i and Linux operating systems. The Flex System p460 enables companies to get the most out of their systems by increasing utilization and performance while reducing costs.

Storage

The IBM Flex System storage capabilities allow you to gain advanced storage-node functionality in your system while taking advantage of your existing storage infrastructure through advanced virtualization. Your Flex System simplifies storage administration with a single storage management user interface that is incorporated into the integrated system manager. This allows you to virtualize third-party storage with non-disruptive migration of the current storage infrastructure.

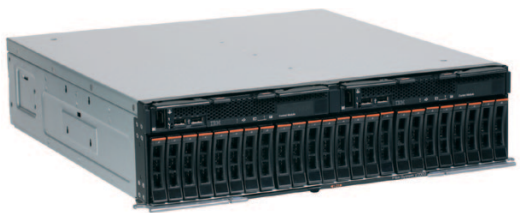

Flex System V7000 Storage Node

The **Flex System V7000 Storage Node** is built on the industry-leading storage virtualization and efficiency capabilities of IBM Storwize® V7000 while being physically and logically integrated into the Flex System. The Flex System V7000 enables external virtualization, consolidation and tiering, and is designed to improve application availability and resource utilization for any size organization. Offering you easy-to-use, efficient and cost-effective management capabilities for new and existing storage resources, Flex System V7000 delivers a no-compromise storage design combined with integrated virtualization, efficiency and performance capabilities. This helps simplify and speed deployment by automating and streamlining provisioning for greater responsiveness and lower costs. And because it can be used in conjunction with the Storwize V7000 system, the Flex System V7000 allows you to progress while maximizing your current storage investments.

The **IBM Flex System Storage Expansion Node** provides dedicated hard disk drives (HDDs) or solid state disk (SSD) storage to an IBM Flex System x220 and x240 compute node, which expands the compute node's capability. Supporting up to 12 SAS/SATA hot-swappable drives, the Flex System Storage Expansion Node is ideal for a variety of application environments, including unstructured data analytics within a distributed database environment and network Attached Storage infrastructure solutions.

Systems management

System administrators face extraordinary demands as they try to plan, document and roll out IT infrastructures, identify capacity needs, get the most out of current assets, and contend with constant budget pressure. More and more the IT infrastructure is being virtualized and administrators must manage complex interactions between virtualization, compute, storage and network platforms. This task is further complicated by multiple tools that don't provide a complete picture of these interlocking infrastructure elements.

IBM Flex System Manager is designed to help you get the most out of your IBM Flex System by automating repetitive tasks, and providing visibility and control across compute, storage, network, and virtualization functions within the data center. With more automation and integrated management across infrastructure elements, IT administrators can focus a greater portion of their time on programs that drive innovation and business advantage.

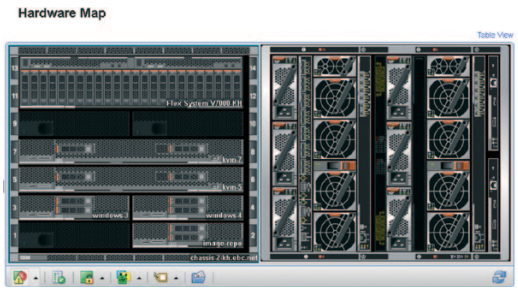

Flex System Manager delivers integrated management across server, storage, network and virtualized resources.

Flex System Manager provides critical management features at each step in the Flex System management lifecycle, including:

- Single point of control: management integration of physical and virtual compute, storage and networking resources from a single management console.
- Smart provisioning: simple and repeatable hardware configuration and bare metal provisioning for server nodes and five click creation and zoning of storage volumes.
- Virtual Appliance Deployment for KVM and PowerVM: deploy virtual appliances that embed the OS and application payload and automatically provision networking and storage resources.
- Upward Integration: integrate hardware management and control functions with VMware vCenter or Microsoft System Center to enable infrastructure management through broadly used virtualization managers.
- Mobile management: monitor your IBM PureFlex and Flex System infrastructure from anywhere using your iOS, Android or Blackberry mobile device.
- Simple administration: Flex System Manager fuel gauge monitors FSM management capacity and makes recommendations for FSM performance optimization.
- Increased scalability: Manage up to 16 Flex System virtualized chassis, 224 compute nodes, and a total of 5000 end points from a single Flex System Manager.

With Flex System Manager, managing your Flex System infrastructure just got a lot easier.

IBM Flex System Enterprise Chassis at a glance	
Height	440 mm (10 EIA rack standard units)
Width	447 mm (EIA 19-inch rack standard width, minus 3 mm clearance)
Depth	800 mm (measured from front bezel to rear of chassis)
	847 mm (measured from ITE latch handle to the power supply handle)
Rack-mount weight	505 lbs
Minimum loadout weight	218 lbs
Maximum loadout weight	493 lbs
Voltage nominal	200 - 240 VAC + 10%
Frequency nominal	50/60Hz

IBM Flex System Enterprise Chassis at a glance	
Power minimum	400 W (estimate)
Power maximum	12,900 watts (12.9 kW)
Maximum Input Current	13.85 A per supply
Node bays	14 standard node bays (7 double wide)
Power supplies§	2/6 2500 W or 2100 W
	200 - 240 VAC
Fabric bandwidth	10 Gb
Switch modules	Optional
80 mm fans	04/08/14
40 mm fans	02/02/14
CMM	½

	x220 Compute Node	x222 Compute Node (per twin)	x240 Compute Node	x440 Compute Node
Form factor	Flex System standard node	Flex System standard node	Flex System standard node	Flex System double-wide node
Processor	2 Intel Xeon E5-2400 Series Processor; 16 cores	2 Intel Xeon E5-2400 Series Processor; 16 cores per twin	2 Intel Xeon E5-2600 v2 Series Processor; 24 cores	4 Intel Xeon E5-4600 Series Processor; 16 cores
Cache	Up to 20 MB per core	Up to 20 MB per core	Up to 20 MB per core	Up to 20 MB per core
Memory	12 DDR3/DDR3L LP, 384 GB max with 32 GB LRDIMMs	12 DDR3/DDR3L LP, 384 GB max per twin with 32 GB LRDIMMs	24 DDR3/DDR3L LP, 768 GB Max with 32 GB LRDIMM	48 DDR3/DDR3L LP, 1.5 TB (32 GB LRDIMMs)
Internal storage	2 x HS 2.5-inch (SAS/SATA/SSD)	1 x 2.5-inch (SATA/SSD), 2 x hot-swap 1.8-inch SSD	2 x HS 2.5-inch (SAS/SATA/SSD)	2 x HS 2.5-inch (SAS/SATA/SSD)
Internal RAID	Software RAID, RAID-0, -1; optional RAID-0, -1, -5, -6, -10, -50	Not applicable	Hardware RAID, RAID-0, -1; optional RAID-0,1, -5, -6, -10, -50	Hardware RAID, RAID-0, -1; optional RAID-0, -1, -5, -6, -10, -50
Internal USB	2 x Standard USB Flash Key + 1 x Front Access USB Key	2 x Standard USB Flash Key + 1 x Front Access USB Key	2 x Standard USB Flash Key + 1 x Front Access USB Key	2 x Standard USB Flash Key + 1 x Front Access USB Key
Ethernet	Dual 1 GbE	IBM Virtual Fabric 2 x 10 GbE LOM support up to 11 x222 nodes with base module using flexible port mapping	Dual 10 GbE	4 x 10 GbE
Chassis support	Flex System Enterprise Chassis	Flex System Enterprise Chassis	Flex System Enterprise Chassis	Flex System Enterprise Chassis
Power management	P-state Capping, Power Maximizer	P-state Capping, Power Maximizer	Active Energy Management	Active Energy Management
Warranty	3 year	3 year	3 year	3 year
Management	IMM2, RTMM KVM Dongle	IMM2, RTMM KVM Dongle	IMM2, RTMM KVM Dongle	IMM2, RTMM KVM Dongle
Operating systems	Microsoft Windows Server, SUSE, Red Hat Enterprise Linux, VMware	Microsoft Windows Server, SUSE, Red Hat Enterprise Linux, VMware	Microsoft Windows Server, SUSE, Red Hat Enterprise Linux, VMware	Microsoft Windows Server, SUSE, Red Hat Enterprise Linux, VMware
RAS features	Chassis redundant/hot-plug power and cooling	Chassis redundant/hot-plug power and cooling	Chassis redundant/hot-plug power and cooling	Chassis redundant/hot-plug power and cooling
	Front panel and FRU/ CRU LEDs	Front panel and FRU/ CRU LEDs	Front panel and FRU/ CRU LEDs	Front panel and FRU/ CRU LEDs

IBM Flex System X6 Compute Node Family at a glance

	Flex System x280 X6	**Flex System x480 X6**	**Flex System x880 X6**
Processors (max)	Intel Xeon E7- 2800 v2 processor families up to 2.8 GHz, up to 1600 MHz memory access, 15 cores per processor	Intel Xeon E7-4800 v2 processor families up to 2.8 GHz, up to 1600 MHz memory access, 15 cores per processor	Intel Xeon E7-8800 v2 processor families up to 3.4 GHz, up to 1600 MHz memory access, 15 cores per processor
Cache (max)	Up to 37.5 MB		
Memory (max)	Up to 3 TB, 48 DIMM slots supporting 64 GB LRDIMMs, RDIMMs	Up to 6 TB, 96 DIMM slots supporting 64 GB LRDIMMs, RDIMMs	Up to 12 TB, 192 DIMM slots supporting 64 GB LRDIMMs, RDIMMs*
Scalability	Two processors max	Scalable to four processors max	Scalable to eight processors max
Ultra-low latency flash storage	Up to 6.4 TB, 16 x 400 GB eXFlash DIMMs	Up to 12.8 TB, 32 x 400 GB eXFlash DIMMs	Up to 12.8 TB, 32 x 400 GB eXFlash DIMMs
Expansion slots	Up to 4 PCIe Gen 3; up to 2 x 16 slots; up to 2 x 8 slots	Up to 8 PCIe Gen 3; up to 4 x 8 slots; up to 4 x 16 slots	Up to 16 PCIe Gen 3; up to 8 x 8 slots; up to 8 x 16 slots
Disk bays (total/hot-swap)	Up to two 2.5-inch Serial Attached SCSI (SAS) hard disk drives (HDDs) or SAS solid state drives (SSDs); or up to eight1.8-inch eXFlash SSDs (4 external, 4 internal)	Up to four 2.5-inch Serial Attached SCSI (SAS) hard disk drives (HDDs) or SAS solid state drives (SSDs); or up to sixteen 1.8-inch eXFlash SSDs (8 external, 8 internal)	Up to eight 2.5-inch Serial Attached SCSI (SAS) hard disk drives (HDDs) or SAS solid state drives (SSDs); or up to thirty-two 1.8-inch eXFlash SSDs (16, external, 16 internal)
Maximum internal storage	Up to 2.4 TB (2 x 2.5-inch SAS/SATA HDDs) or up to 3.2 TB (2 x 2.5-inch SSDs) or 3.2 TB (8 x 1.8-inch eXFlash SSDs)	Up to 4.8 TB (4 x 2.5-inch SAS/SATA HDDs) or up to 6.4 TB (4 x 2.5-inch SSDs) or 6.4 TB (16 x 1.8-inch eXFlash SSDs)	Up to 9.6 TB (8 x 2.5-inch SAS/ SATA HDDs) or up to 12.8 TB (8 x 2.5-inch SSDs) or 12.8 TB (32 x 1.8-inch eXFlash SSDs)
Network interface	Integrated 4 x10 GbE ports; supports 8/16 Gb FC and QDR IB	Integrated 8 x10 GbE; supports 8/16 Gb FC and QDR IB	Integrated 16 x10 GbE; supports 8/16 Gb FC and QDR IB
Hot-swap components	HDDs, SSDs		
RAID support	RAID-0, RAID-1; optional RAID-5		
Systems management	IMM2, next-generation IBM light path diagnostics, IBM FSM Director with Scaling Support		
Operating systems supported	Microsoft Windows Server, Red Hat Enterprise Linux Server, SUSE Linux Enterprise Server, VMware vSphere Hypervisor		
Limited warranty	3-year customer replaceable unit and onsite service, next business day 9 × 5, service upgrades available		

	p260 Compute Node	p270 Compute Node	p460 Compute Node
Form factor	Flex System standard node	Flex System standard node	Flex System double-wide node
Processor cores	4, 8 or 16 cores, POWER7+, 64-bit processors with VSX, Memory Expansion acceleration and Encryption acceleration Configuration Options: 2-core 4.0 GHz, 4-core 4.0 GHz, 8-core 3.6 GHz, 8-core 4.1 GHz	24 cores, POWER7+, 64-bit processors with VSX, Memory Expansion acceleration and Encryption acceleration Configuration options: 24-core 3.1 GHz or 3.4 GHz	16 or 32 cores, POWER7 64-bit processors with AltiVec SIMD and Hardware Decimal Floating-Point acceleration or 8 or 16 cores, POWER7+, 64-bit processors with VSX, Memory Expansion acceleration and Encryption acceleration. Configuration Options: 4-core 3.3 GHz or 4.0 GHz, 8-core 3.2 GHz or 3.6 GHz, 8-core 3.5 GHz or 4.1 GHz
Cache	Level 2 (L2): 256 KB per processor core Level 3 (L3): 10 MB per processor core on POWER7+ offerings	Level 2 (L2): 256 KB per processor core Level 3 (L3): 10 MB per processor core	Level 2 (L2): 256 KB per processor core Level 3 (L3): 4 MB per processor core on POWER7 10 MB per processor core on POWER7+ offerings
Memory (min/max)	8 GB up to 512 GB, 16 DIMM slots, ECC IBM Chipkill DDR3 SDRAM running at 1066 MHz	8 GB up to 512 GB, 16 DIMM slots, ECC IBM Chipkill DDR3 SDRAM running at 1066 MHz plus Active Memory Expansion with hardware assist	16 GB up to 1 TB node, 32 DIMM slots, ECC IBM Chipkill DDR3 SDRAM running at 1066 MHz
Internal disk storage	Up to two 2.5-inch HDDs or two 1.8-inch SSDs	Up to two 2.5-inch HDDs or two 1.8-inch SSDs	Up to two 2.5-inch HDDs or two 1.8-inch SSDs
Networking/ Expansion	Two PCIe expansion slots	Two PCIe expansion slots One ETE adapter card slot (provides dedicated Dual VIOS on internal drives with optional adapter)	Four PCIe expansion slots
Systems management	Integrated systems management processor, light path diagnostics, Predictive Failure Analysis, CSM, Serial Over LAN, IPMI compliant	Integrated systems management processor, light path diagnostics, Predictive Failure Analysis, CSM, Serial Over LAN, IPMI compliant	Integrated systems management processor, light path diagnostics, Predictive Failure Analysis, CSM, Serial Over LAN, IPMI compliant
Operating systems	AIX 6.1 and AIX 7.1 IBM i 6.1 and 7.1 RHEL 5.7, 6.2 SLES11 SP2	AIX 6.1 and AIX 7.1 IBM i 6.1 and IBM i 7.1 RHEL 6.4 SLES11 SP2	AIX 6.1 and AIX 7.1 IBM i 6.1 and 7.1 RHEL 5.7, 6.2 SLES11 SP2
RAS features	Chassis redundant/hot-plug power and cooling Front panel and FRU/CRU LEDs Concurrent code update Processor deallocation compute node hot-plug Dual AC Power Supply Auto reboot on power loss Internal and chassis-external temperature monitors 64B Marking ECC code supporting x8 IS DDR3 DIMMs System mgmt alerts IBM Chipkill ECC detection and correction	Chassis redundant/hot-plug power and cooling Front panel and FRU/CRU LEDs Concurrent code update Processor deallocation compute node hot-plug Dual AC power supply Auto reboot on power loss Internal and chassis-external temperature monitors 64B Marking ECC code supporting x8 IS DDR3 DIMMs System mgmt alerts IBM Chipkill ECC detection and correction	Chassis redundant/hot-plug power and cooling Front panel and FRU/CRU LEDs Concurrent code update Processor deallocation compute node hot-plug Dual AC Power Supply Auto reboot on power loss Internal and chassis-external temperature monitors 64B Marking ECC code supporting x8 IS DDR3 DIMMs System mgmt alerts IBM Chipkill ECC detection and correction
Energy management	EnergyScale energy management	EnergyScale energy management	EnergyScale energy management

	IBM Flex System PCIe Expansion Node
Compute nodes supported	Flex System x220 and x240 compute nodes
Supported bus widths	1x, 2x, 4x, 8x and 16x*
Adapter capacity	Two full-length, full-height x16 Two half-length, half-height x8
Adapter sizes	Standard-height adapters, 4.20-inch (106.7 mm) Low-profile adapters, 2.536-inch (64.4 mm) Half-length adapters, 6.6-inch (167.65 mm) Full-length adapters, 12.283-inch (312 mm)
Adapter quantities	Up to two low-profile adapters Up to two full-height adapters Up to one full-height, doublewide adapters
PCIe standards	Supports 1.1 and 2.0
Generation	Generation 1- and Generation 2-compliant adapters†
Power	Supports greater than 75 W PCIe adapters using standard PCIe auxiliary power connectors‡
Form factor	Flex System standard node
Chassis support	Flex System Enterprise Chassis
I/O expansion	2x Mezz adapters (x8+x4)
Warranty	3-year customer replaceable unit and onsite service, next business day 9×5, service upgrades available
Management	IMM2, RTMM KVM Dongle

IBM Flex System V7000 Storage Node at a glance	
Host interface	SAN-attached 8 Gbps Fibre Channel, 10 Gigabit Ethernet (GbE) Fibre Channel over Ethernet and iSCSI host connectivity
Cache per controller/control enclosure/clustered system	8 GB/16 GB/64 GB
Supported drives	2.5-inch disk drives • 500 GB, 1 TB 7,200 rpm nearline SAS • 14 6 and 300 GB 15,000 rpm SAS • 300, 600 and 900 GB 10,000 rpm SAS 200 and 400 GB E-MLC SSD
Maximum drives supported	240 per control enclosure; 960 per clustered system
Rack support	IBM Flex System Enterprise Chassis
Management software	IBM Flex System Manager
Advanced features included with each system	Easy Tier, IBM FlashCopy®, internal virtualization and thin provisioning, data migration, system clustering
Optional features	Remote mirroring, Real-time Compression, external virtualization

IBM Flex System Storage Expansion Node at a glance	
Compute nodes supported	Flex System x220 and x240 compute nodes
Feature on Demand	ServeRAID M5100 Series RAID-6 Upgrade ServeRAID M5100 Series SSD Caching Enabler ServeRAID M5100 Series Performance Accelerator
HDD support	SAS 2.1 at 6 Gbps and SATA
PCIe support	x8 PCIe Gen 3, 8 GTps
Drive modes	JBOD and RAID
RAID modes	RAID-0, -1, -5, -6, -10, -50 and -60
Cache options	512 MB and 1 GB with cache-to-flash super capacitor offload
Form factor	Flex System standard node
Chassis support	Flex System Enterprise Chassis
Limited warranty	3-year customer replaceable unit and onsite limited warranty, next business day 9×5, service upgrades available
Systems management	IMM2, RTMM KVM Dongle

	IBM Flex System EN4132 2-port 10 Gb Ethernet Adapter	**IBM Flex System EN2024 4-port 1 Gb Ethernet Adapter**
Use	Supports RDMA and RoCE technologies, helping provide application acceleration and low latency for specialized applications. This adapter will work with the 10 Gb Flex System Fabric Switch and the 10 Gb Ethernet Pass-Thru Module.	Supports IO virtualization features like VMware, NetQueue and Microsoft VMQ technologies. When combined with the Flex System EN2092 1Gb Ethernet Scalable Switch, you get an end-to-end 1 Gb solution in Flex System chassis.
Ports	Two 10 Gb Ethernet ports	Four 1 Gb ports
Technology	10 Gb Ethernet based on Mellanox Connect X3 ASIC	1 Gb Ethernet

	IBM Flex System Fabric EN4093R 10 Gb Scalable Switch	**IBM Flex System Fabric SI4093 System Interconnect Module**	**IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch**
Use	Supports 10 Gb Ethernet - offering scalability, flexibility and performance to help address many networking concerns today while providing capabilities that will better prepare clients for the future.	Preconfigured transparent device designed for easy connectivity, reduced management and setup time, while providing the scalability, flexibility and performance clients need now and in the future. Element of the IBM Flex System Interconnect Fabric solution	Unsurpassed convergence flexibility, scalability and performance, while delivering networking innovations to help clients address today's converged networking requirement and their potential future needs. The switch is designed to simplify network resource provisioning and resource optimization based on application requirements.
Ports	Up to 42 internal 10 Gb ports and 22 external 10 Gb ports. Eight of these 10 Gb uplinks can be used as two 40 Gb uplink ports. Flexible port mapping offers dynamic configuration of active ports	Up to 42 internal 10 Gb ports and 22 external 10 Gb ports. Eight of these 10 Gb uplinks can be used as two 40 Gb uplink ports. Flexible port mapping offers dynamic configuration of active ports	Up to 42 internal ports, and 16 external ports (2 x 10 Gb SFP+, 12 x OmniPorts, 2 x 40 Gb QSFP+). Flexible port mapping offers dynamic configuration of active ports
Technology	40 Gb, 10 Gb and 1 Gb Ethernet Includes stackable FCoE transit switch capability	Preconfigured loop-free interface that does not rely upon spanning tree protocols;40 Gb, 10 Gb and 1 Gb Ethernet; Includes FCoE transit switch capability	4 Gb and 8 Gb Fibre Channel 40 Gb, 10 Gb and 1 Gb Ethernet

	IBM Flex System EN2092 1 Gb Ethernet Scalable Switch	**IBM Flex System EN4091 10 Gb Ethernet Pass-Thru Module**	**IBM Flex System CN4054 10 Gb Virtual Fabric**
Use	This 1 Gb Ethernet Scalable switch provides outstanding flexibility allowing you to buy one switch today and enhance its functionality in the future.	Get simple connectivity of the Flex System Chassis to any external network infrastructure.	Supports multiple advanced protocols for Intel processor-based IBM Flex System compute nodes. It can run in a physical, virtual or switch independent modes, and helps provide a way to reduce data center costs by using a common infrastructure for Ethernet and storage networks.
Ports	Up to 28 internal 1 Gb ports and 20 external 1 Gb ports. 20 of these 1 Gb ports can be used as two 10 Gb external port. Flexible port mapping offers dynamic configuration of active ports	14 internal 10 Gb links 14 external 10 Gb SFP+ uplinks	Four 10 Gb ports, each port supports up to four virtual ports, upgrade to run FCoE or HW iSCSI
Technology	1 Gb and 10 Gb Ethernet	1 Gb and 10 Gb Ethernet, 10 Gb Fibre Channel over Ethernet	1 Gb and 10 Gb Ethernet, 10 Gb Ethernet, FCoE and HW iSCSI, based on Emulex BE3 ASIC

	IBM Flex System Fabric EN6131 40 Gb Ethernet Switch	**Cisco Nexus B22 Fabric Extender for IBM Flex System**	**IBM Flex System EN4023 10 Gb Scalable Switch**
Use	Provides end-to-end 40 Gb Ethernet connectivity for IBM Flex System chassis. Features high-performance 40 Gb Ethernet that can also auto negotiate to 10 GbE speed. Enables less than 0.7usec latency node to node, ideal for clients running Ethernet infrastructure in high-speed trading, telecommunications, Web applications, virtualization and Cloud computing.	Provides an extension of the Cisco Nexus fabric to IBM Flex System chassis. Aggregates network connections of all IBM Flex System nodes to upstream Cisco Nexus Switch. Adds new options for Cisco-centric networking organizations.	Streamlines network deployment and operation using Brocade Virtual Cluster Switching (VCS) in a Brocade VDX environment. Available Feature on Demand (FoD) upgrades allow purchase of additional port licenses as needed. User may assign available ports for internal or external connections up to maximum of 42 internal ports and up to 16 external ports.
Ports	14 internal ports 18 external	14 internal 1 Gb/10 Gb Ethernet/Fibre Channel over Ethernet (FCoE) ports 8 external SFP+ 10 Gb ports	24 10 Gb Ethernet ports in base switch External SFP+ 10 Gb Ethernet ports Upgrade 1: sixteen 10 Gb Ethernet ports and two 40 Gb QSP+ ports Upgrade 2: sixteen 10 Gb Ethernet ports
Technology	Mellanox 40 Gb Ethernet	1 Gb and 10 Gb Ethernet 10 Gb Ethernet and FCoE Cisco Nexus Technology Copper breakout cables	10 Gb Ethernet Brocade VCS

For more information about additional Flex System Ethernet Networking options, click here

	IBM Flex System FC5022 8/16 Gb SAN Scalable Switch	IBM Flex System FC3171 8 Gb SAN Switch	IBM Flex System FC3171 8 Gb SAN Pass-Thru Module
Use	Market-leading 16 Gbps Fibre Channel technology as well as expert optimized, automated and integrated capabilities. The switch is designed to support highly virtualized computing and SAN environments with high performance, reliability and usability.	Provides an integrated, simple connection to existing SAN fabrics and storage. Based on QLogic's proven Fibre Channel expertise. The switch is designed to set up quickly and be easy to manage. Minimize time and risk and support faster access to your data faster and quicker and better business decisions.	Enables 8 Gb connectivity to storage from the Flex System chassis and offers enhanced Fibre Channel functions like network port aggregation, auto stream guard, Enhanced N_Port ID Virtualization (NPIV) and automatic failover.
Ports	Up to 48 total physical ports: 28 internal, 20 external. 48 virtual channels per port	20 ports (14 internal, 6 external SFP+)	14 internal 8 Gb ports and six external 8 Gb ports. Works at 4 Gb and 8 Gb speed
Technology	Brocade	QLogic	QLogic

	IBM Flex System FC5022 2-port 16 Gb Fibre Channel Adapter	IBM Flex System FC3052 2-port 8 Gb Fibre Channel Adapter	IBM Flex System FC3172 2-port 8 Gb Fibre Channel Adapter
Use	Enables high-speed access for Flex System compute nodes to external Storage area network (SAN). Offers end-to-end 16 Gb connectivity to your SAN. It can auto-negotiate, and also work at 8 Gb and 4 Gb speeds. It has enhanced features like N-port trunking, as well as increased encryption for security.	Works with any of the 8 Gb or 16 Gb Flex System Fibre Channel switch modules. When compared to the previous-generation 4 Gb adapters, the new generation 8 Gb adapters double throughput speeds for Fibre Channel traffic. As a result, it is possible to manage increased amounts of data.	Works with any of the 8 Gb or 16 Gb Flex System Fibre Channel switch modules. When compared to the previous generation of 4 Gb adapters, the new 8 Gb adapters double the throughput speeds for Fibre Channel traffic. As a result it is possible to manage increased amounts of data.
Ports	Two 16 Gb Fibre Channel ports	Two 8 Gb Fibre Channel ports	Two 8 Gb ports
Technology	Brocade	Emulex	QLogic

For more information about additional Flex System Fibre Channel options, click here

	IBM Flex System IB6131 InfiniBand Switch
Use	Designed to offer the performance you need to support clustered databases, parallel processing, transactional services and high-performance embedded I/O applications, reducing task completion time and lowering cost per operation. Virtual Protocol Interconnect also simplifies system development by serving multiple fabrics with one hardware design.
Ports	14 internal ports
Technology	Mellanox

	IBM Flex System IB6132 2-port FDR InfiniBand Adapter
Use	Designed to meet your critical performance needs. Supports switch-embedded subnet managers and host-based subnet managers.
Ports	2 FDR ports for 56 Gbps bandwidth
Technology	Mellanox

For more information about additional Flex System Infiniband options, click here

Why IBM?

IBM has taken knowledge, expertise and technology gained from decades of experience and investment in IT solutions for business problems and incorporated it into the IBM Flex System. With a commitment to open standards you can integrate IBM solutions with other elements of your own environment with your network of partners, customers and suppliers.

With a broad ecosystem of partners with technical and industry expertise, and the unique ability and skill to integrate it all together for you—along with an unwavering commitment to your success—you can rely on IBM and your IBM Flex System.

For more information

To learn more about the IBM Flex System visit: **ibm.com**/flex or contact your IBM representative or IBM Business Partner.

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. We'll partner with credit-qualified clients to customize an IT financing solution to suit your business goals, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: **ibm.com**/financing

[4] Comparisons are top bin 5600 to top bin 2600 two-socket configuration. For more details, visit www.intel.com/content/www/us/en/benchmarks/workstation/xeon-e5-2600.html

[5] Performance comparison using SPECint*_rate_base2006 benchmark result divided by the processor TDP. For more details, visit www.spec.org/cpu2005/results/res2011q4/cpu2006-20111121-19037.html

[6] 100 percent performance improvement is based on preliminary results of SPECint*_rate_base2006 and SPECfp*_rate_base2006, plus performance gains from eXFlash DIMM storage. SPEC benchmark results will be available after 4/8/14. Configurations: 4-socket x480 X6 server using Intel Xeon processor E7-4890 v2 vs. 4-socket server using the previous top-of-the-line E7-4870 (v1).

[7] Based on a 4-socket IBM Flex System x480 X6 or 8-socket x880 X6 configuration with the capability of supporting 32 Flash DIMM slots x 400 GB Flash DIMMs.

IBM Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
April 2014

* PCIe mechanical connectors limit larger lane add-in cards from being inserted into smaller mechanical connectors.

† Generation 3-compliant adapters will operate using Generation 2 compliance.

‡ Contact your IBM representative for additional information on other high-wattage adapters that can be supported.

§ 2100 W power supplies are also available as an option for select compute node configurations.

[1] Based on double-dense Flex System x222 compute nodes with 7:1 oversubscription. For more details, visit **ibm.com**/systems/flex/networking/ethernet.html

[2] IBM Flex System Flash on the Flex System x240 compute node supports up to 1.6 TB compared to the Dell M610 with 800 GB.

[3] Manage one IBM RackSwitch G8264CS versus 20 devices (two G8264CS switches and 18 SI4093 interconnect modules).

Please Recycle

WZB12345-USEN-05