

Flex System Enterprise Chassis

Lenovo Press Product Guide

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

The Flex System Enterprise Chassis is shown in Figure 1.



Figure 1. The Flex System Enterprise Chassis

Did you know?

Flex System is the platform for PureFlex System. PureFlex System combines advanced hardware and software along with patterns of expertise and integrates them into three optimized configurations, Express, Standard, and Enterprise, that are simple to acquire and deploy so you get fast time to value for your solution.

Key features

The Flex System Enterprise Chassis is a simple, integrated infrastructure platform that supports a mix of compute, storage, and networking resources to meet the demands of your applications. The solution is easily scalable with the addition of another chassis with the required nodes. With Flex System Manager, multiple chassis can be monitored from a single screen. The 14 node, 10U chassis delivers high-speed performance complete with integrated servers, storage, and networking. This flexible chassis is designed for a simple deployment now and to scale to meet your needs in the future.

Flexibility and efficiency

The 14 bays in the chassis allow the installation of compute or management nodes, with networking modules in the rear. A single chassis or a group of chassis can be fully customized to the specific needs of the computing environment. With support for POWER7® and Intel processor-based nodes, you can choose the architecture you need. IT can meet the needs of the business using a single system across multiple architectures and operating environments.

The system monitors and manages power usage on all major chassis components so you have total control over power consumption. Available power supply options are AC, -48V DC and HVDC. AC power options can be configurable in either a single or three-phase power domain. The chassis supports N+N or N+1 redundant power supplies and an entirely passive mid-plane to meet your reliability needs. The power supplies are 80 PLUS Platinum-certified indicating high energy efficiency. The chassis design also optimizes cooling with cooling zones within the chassis. The system manages the fan modules based on node configuration within the chassis. So, the system can increase the speed of certain fan modules to cool potential hot spots, and use lower speeds for other fan modules where appropriate.

Easily scalable with simple administration

Because the Flex System Enterprise Chassis is an all-in-one solution, it is designed for growth from a single chassis to many. Adding compute, storage, or networking capability is as simple as adding additional nodes, modules, or chassis. The simple, highly integrated management system allows you to use the Chassis Management Modules integrated into each chassis to administer a single chassis, or Flex System Manager controls up to 16 chassis from a single panel.

Designed for multiple generations of technology

The Flex System Enterprise Chassis is designed to be the foundation of your IT infrastructure now and into the future. Compute performance requirements are always on the rise and networking demands continue to grow with rising bandwidth needs and a shrinking tolerance for latency. The chassis is designed to scale to meet the needs of your future workloads, offering the flexibility to support current and future innovations in compute, storage, and networking technology.

Locations of key components and connectors

Figure 2 shows the front of the Enterprise Chassis.




Figure 2. Front of the Flex System Enterprise Chassis

Figure 3 shows the rear of the Enterprise Chassis.




Figure 3. Rear of the Flex System Enterprise Chassis

Standard specifications

The following table lists the standard specifications.

Table 1. Standard specifications

Components	Specification
Model	System x® ordering sales channel: machine type 8721 Power Systems sales channel: MTM 7893-92X
Form factor	10U rack-mounted unit.
Maximum number of compute nodes	14 half-wide (single bay), seven full-wide (two bays), three double-height full-wide (four bays) or one quadruple-height full-wide (8 bays). Mixing is supported.
Chassis per 42U rack	Four.
Nodes per 42U rack	56 half-wide, or 28 full-wide.
Management	One or two Chassis Management Modules (CMMs), for basic chassis management. Two CMMs form a redundant pair; One CMM standard in 8721-A1x. The CMM interfaces with the integrated management module (IMM) or Flexible Service Processor (FSP) integrated in each compute node in the chassis. Optional Flex System Manager management appliance for comprehensive management, including virtualization, networking, and storage management.
I/O architecture	Up to eight lanes of I/O to an I/O adapter card, with each lane capable of up to 16 Gbps bandwidth. Up to 16 lanes of I/O to a half-wide node with two adapters. A wide variety of networking solutions, including Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and InfiniBand.
Power supplies	Up to six power supplies that can provide N+N or N+1 redundant power. Power supplies are 80 PLUS Platinum certified and provide over 94% efficiency at 50% load and 20% load. Each power supply contains two independently powered 40 mm cooling fan modules. Power supply options are 2500W AC, 2100W AC, 2500W HVDC and 2500W -48V DC power supplies. All installed power supplies must be identical.
Fan modules	Ten fan modules (eight 80 mm fan modules and two 40 mm fan modules); Four 80 mm and two 40 mm fan modules standard.
Dimensions	Height: 440 mm (17.3 in) Width: 447 mm (17.6 in) Depth: 800 mm (31.5 in) (measured from front bezel to rear of chassis), 840 mm (33.1 in) (measured from node latch handle to the power supply handle)
Weight	Minimum configuration: 96.62 kg (213 lb) Maximum configuration: 220.45 kg (486 lb)
Sound level	7.5 bels (declared)
Temperature	Operating air temperature 5°C to 40°C
Electrical power	AC Input power: 200 - 240 V ac (nominal), 50 or 60 Hz HVDC DC Input power: 240 - 380 V dc (nominal) -48 V DC Input power: -48V to -60 V dc (nominal) Minimum configuration: 0.51 kVA (two power supplies) Maximum configuration: 13 kVA (six power supplies)
Power consumption	12,900 watts maximum

The following table lists what each standard model, Express model and TopSeller model includes. In addition, each model ships with:

- One C19 to C20 two-meter power cable for each power supply
- One Rack Mount Kit

Models

The following table lists the specifications of the standard Enterprise Chassis models.

Table 2. Models

Model	Node bays	CMM (2 max)	FSM (1 max)	I/O bays (used / max)	I/O modules included	Power supplies (6 max)	40mm fans (2 max)	80mm fan (8 max = 4 pairs)	Console breakout cable
Standard models									
8721-A1x	14	1	Optional	0 / 4	None	2x 2500W	2	4	1
8721-LRx	14	1	Optional	0 / 4	None	2x 2100W	2	4	1
8721-DCx	14	1	Optional	0 / 4	None	2x 2500W -48V DC	2	4	1
Express models - Europe									
8721-K1G	14	1	Optional	0 / 4	None	2x 2500W	2	4	1
TopSeller models - North America									
8721-E1Y	14	2	Optional	2 / 4	2x SI4093	6x 2500W	2	8	1
8721-E2Y	14	2	Optional	2 / 4	2x CN4093	6x 2500W	2	8	1

Supported compute nodes

The following table lists the compute nodes that are supported in the Flex System Enterprise Chassis. The table also lists the maximum number installable.

Table 3. Supported compute nodes and maximum quantities

Description	Machine type	Maximum number of nodes per chassis*
Flex System x220 Compute Node	7906	14
Flex System x222 Compute Node	7916	14
Flex System x240 Compute Node	8737	14
Flex System x240 M5 Compute Node	9532	14
Flex System x440 Compute Node	7917	7
Flex System x440 Compute Node (E5-4600 v2)	7167	7
Flex System x280 X6 Compute Node	7903	7
Flex System x480 X6 Compute Node	7903	7
Flex System x880 X6 Compute Node	7903	7
Flex System p24L Compute Node	1457	14
Flex System p260 Compute Node	7895-22X	14
Flex System p270 Compute Node	7954-24X	14
Flex System p460 Compute Node	7895-42X	7
Flex System Manager Node	8731-A1x	1 (2 per 16 chassis)

* The actual number of compute nodes systems that can be powered on in a chassis depends on these factors:

- The TDP power rating for the processors that are installed in the compute nodes (x86 servers)
- The number of power supplies installed in the chassis
- The capacity of the power supplies installed in the chassis (2100 W or 2500 W)
- The power redundancy policy used in the chassis (N+1 or N+N)

For the x480 X6 or x880 X6, if two compute nodes are scaled together, then a maximum of 3 scaled complexes can be installed. For the x880 X6, if four compute nodes are scaled together, then a maximum of 1 scaled x880 X6 complexes can be installed. Other bays can be filled with other compute nodes.

The table in the Power Supplies section provides guidelines about what number of compute nodes can be powered on in the Flex System Enterprise Chassis, based on the type and number of power supplies installed.

See IBM® ServerProven® at the following web address for the latest information about the servers supported in the Flex System Enterprise Chassis:

<http://ibm.com/servers/eserver/serverproven/compat/us/>

Supported I/O modules

The Flex System Enterprise Chassis has four high-speed switch bays that are capable of supporting a variety of I/O architectures.

The switches are installed in switch bays in the rear of the Flex System Enterprise Chassis as shown in the following figure. Switches are normally installed in pairs (bays 1 & 2, and bays 3 & 4), because I/O adapter cards installed in the compute nodes route to two switch bays for performance and redundancy.




Figure 4. Location of the switch bays in the Flex System Enterprise Chassis

The following figure shows one of the available switches, the Flex System Fabric EN4093 10Gb Scalable Switch. It offers the following interface ports:

- Up to 42 internal 10Gb ports, three to each half-wide compute node bay (not shown in the figure)
- Up to 14 external 10Gb ports with SFP+ connectors (license upgrades enable all ports)
- Up to 2 external 40Gb uplink ports with QSFP+ connectors (license upgrades enable these ports)



Figure 5. Flex System Fabric EN4093 10Gb Scalable Switch

The following table lists the switches that are available for the chassis.

Table 4. I/O modules and upgrades

Description	Part number	System x feature code	Power Systems feature code (7893-92X)
Ethernet switches			
Flex System EN2092 1Gb Ethernet Scalable Switch	49Y4294	A0TF	3598
Flex System EN4091 10Gb Ethernet Pass-thru	88Y6043	A1QV	3700
Flex System Fabric CN4093 10Gb Converged Scalable Switch	00D5823	A3HH	ESW2
Flex System Fabric EN4093R 10Gb Scalable Switch	95Y3309	A3J6	ESW7
Flex System Fabric EN4093 10Gb Scalable Switch	49Y4270	A0TB	3593
Flex System Fabric SI4093 System Interconnect Module	95Y3313	A45T	ESWA
Cisco Nexus B22 Fabric Extender for Flex System	94Y5350	ESWB	ESWB
Flex System EN4023 10Gb Scalable Switch	94Y5212	ESWD	ESWD
Flex System EN6131 40Gb Ethernet Switch	90Y9346	A3HJ	ESW6
Fibre Channel switches			
Flex System FC5022 16Gb SAN Scalable Switch	88Y6374	A1EH	3770
Flex System FC5022 24-port 16Gb SAN Scalable Switch	00Y3324	A3DP	ESW5
Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch	90Y9356	A1EJ	3771
Flex System FC3171 8Gb SAN Switch	69Y1930	A0TD	3595
Flex System FC3171 8Gb SAN Pass-thru	69Y1934	A0TJ	3591
InfiniBand switches			
Flex System IB6131 InfiniBand Switch (QDR/FDR)	90Y3450	A1EK	3699

I/O architecture

Each half-wide compute node (such as the Flex System x240 Compute Node) has two adapter slots, and each full-wide compute node (such as the Flex System x440 Compute Node) has four adapter slots. The adapter slots in each compute node route through the chassis midplane to the switch bays. The architecture supports up to eight ports per adapter although currently only two-port and four-port adapters are available.

The following figure shows how two-port adapters are connected to switches installed in the chassis.

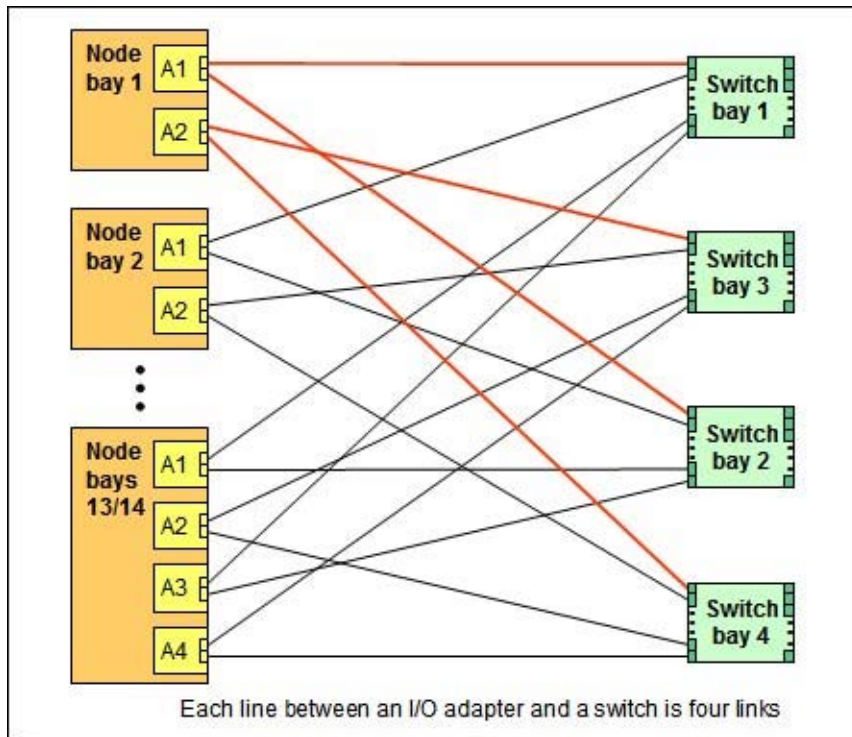


Figure 6. Logical layout of the interconnects between two-port I/O adapters and I/O modules

A four-port adapter doubles the connections between each adapter and switch pair (for example, a four-port adapter in A1 in each compute node routes two connections to switch 1 and two connections to switch 2).

The following figure shows how four-port adapters are connected to switches that are installed in the chassis.

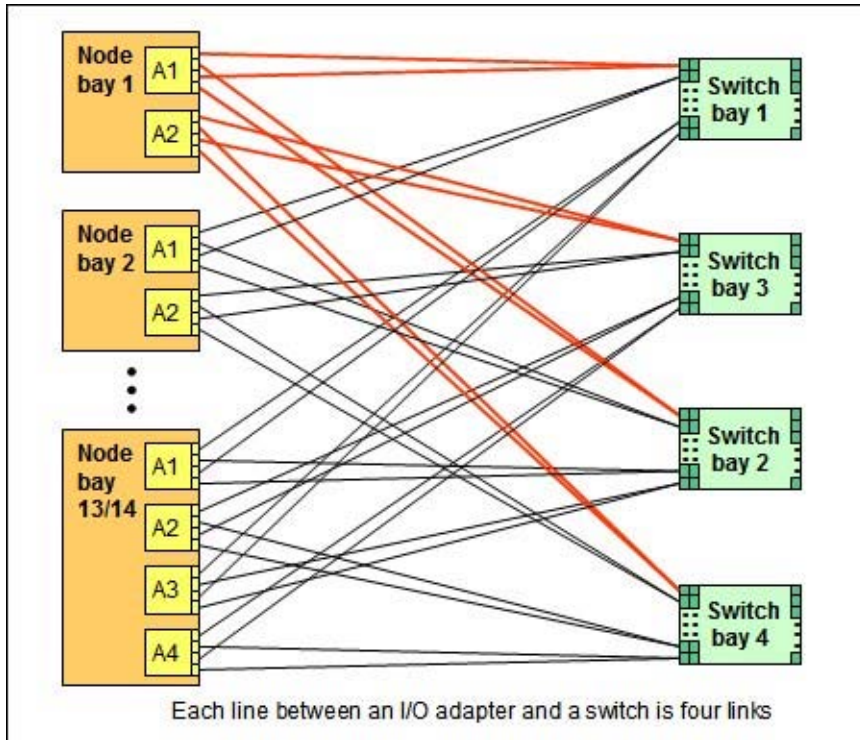


Figure 7. Logical layout of the inter-connections between four-port I/O adapters and I/O modules

To make use of all four ports of a four-port card, you must install two switches in the chassis, and each switch must have 28 internal ports enabled. The Flex System Fabric EN4093 10Gb Scalable Switch with license upgrade 1 enables 28 internal ports, for example. Similarly, when 6-port cards become available, two switches must be installed, and 42 internal ports must be enabled. The Flex System Fabric EN4093 10Gb Scalable Switch with license upgrade 2 enables a total of 42 internal ports, for example.

The following table shows the connections between adapter slots in the compute nodes to the switch bays in the chassis.

Table 5. Adapter to I/O bay correspondence

I/O adapter slot in the server	Port on the adapter	Corresponding I/O module bay in the chassis
Slot 1	Port 1	Module bay 1
	Port 2	Module bay 2
	Port 3 (for 4-port cards)*	Module bay 1
	Port 4 (for 4-port cards)*	Module bay 2
Slot 2	Port 1	Module bay 3
	Port 2	Module bay 4
	Port 3 (for 4-port cards)*	Module bay 3
	Port 4 (for 4-port cards)*	Module bay 4
Slot 3 (full-wide compute nodes only)	Port 1	Module bay 1
	Port 2	Module bay 2
	Port 3 (for 4-port cards)*	Module bay 1
	Port 4 (for 4-port cards)*	Module bay 2
Slot 4 (full-wide compute nodes only)	Port 1	Module bay 3
	Port 2	Module bay 4
	Port 3 (for 4-port cards)*	Module bay 3
	Port 4 (for 4-port cards)*	Module bay 4

* To make use of all four ports of a four-port adapter, the switch must have 28 internal ports enabled, and two switches must be installed in the bays as indicated.

Chassis Management Module

The Chassis Management Module (CMM) provides single-chassis management. The CMM is used to communicate with the management controller in each compute node (IMMv2 in Intel processor-based compute nodes and FSP in POWER7® processor-based compute nodes) to provide system monitoring, event recording and alerts, and to manage the chassis, its devices, and the compute nodes.

The chassis has one CMM installed standard but supports two CMMs for redundancy. If one CMM fails, the second CMM can detect its inactivity and activate itself to take control of the system without any disruption. The CMM is central to the management of the chassis and is required in the Enterprise Chassis.

The following table shows the ordering information.

Table 6. Chassis Management Module

Part number	Feature codes first / second	Description	Standard / Maximum
68Y7030	A0TM / A0UE	Flex System Chassis Management Module	1 / 2

The following figure shows the Chassis Management Module. See Figure 3 for the location of the CMM in the chassis.




Figure 8. Chassis Management Module

The CMM provides these functions:

- Power control
- Fan management
- Chassis and compute node initialization
- Switch management
- Diagnostics: chassis, I/O options, and compute nodes
- Resource discovery and inventory management
- Resource alerts and monitoring management
- Chassis and compute node power management
- Security policy management
- Role-based access control
- Support for up to 84 local CMM user accounts
- Support for up to 32 simultaneous sessions

The CMM has the following connectors:

- USB connection. This connection can be used for insertion of a USB media key for tasks, such as firmware updates.
- 10/100/1000 Mbps RJ45 Ethernet connection to connect to a management network. The CMM can be managed via this Ethernet port.
- Serial port (mini-USB) for local command-line interface (CLI) management. Use serial cable 90Y9338 for connectivity.

The CMM has the following light-emitting diodes (LEDs) that provide the following information:

- Power-on LED
- Activity LED
- Error LED
- Ethernet port link and port activity LEDs

The CMM also incorporates a reset button, which, when pressed, resets the CMM back to its default condition. It has two functions, depending on how long the button is pressed:

- When pressed for less than 5 seconds, the CMM restarts.
- When pressed for more than 5 seconds (for example, 10 or 15 seconds), the CMM configuration is reset to the manufacturing defaults, and then the CMM restarts.

The CMM supports a web-based graphical user interface (GUI) that provides a way to perform CMM functions within a supported web browser. You can also perform management functions through the CMM command-line interface (CLI). Both the web-based GUI and the CLI are accessible via the single RJ45 Ethernet connector on the CMM or from any other system that is connected to the same (management) network.

The CMM has the following default static IPv4 address. By default, the CMM is configured to respond to Dynamic Host Configuration Protocol (DHCP) first before using its static IPv4 address:

- IP address: 192.168.70.100
- Subnet: 255.255.255.0
- User ID: USERID (all capital letters)
- Password: PASSWORD (all capital letters, with a zero instead of the letter O)

The CMM does not have a fixed static IPv6 IP address, by default. Initial access to the CMM in an IPv6 environment can be performed by either using the IPv4 IP address or the IPv6 link-local address. The IPv6 link-local address is automatically generated based on the Media Access Control (MAC) address of the CMM.

The CMM is the key component enabling the integrated management network. Internally, the CMM has a multiple port L2 1Gigabit Ethernet switch with dedicated links to all 14 node bays, all four switch bays, and the second CMM, if installed. These connections are all point-to-point, ensuring dedicated bandwidth. The 1GbE links are full-duplex, fixed speed (not auto-negotiate) links. The 1 GbE management network is only accessible by each node's management controller (IMMv2 or FSP), each switch module's management interfaces, the CMM, and the Flex System Manager (FSM) management appliance. This design permits the separation of the management network from the data network.

The CMM has a high-security policy that is enabled by default, which means that the following policies are enabled by default:

- Strong password policies with automatic validation and verification checks
- Required update of the default passwords after the initial setup
- Only secure communication protocols, such as SSH and SSL. Unencrypted protocols, such as HTTP, Telnet, and SNMPv1, are disabled.
- Certificates to establish secure, trusted connections for applications that run on the management processors

Flex System Manager

Flex System Manager (FSM) is a systems management solution that offers flexible integrated systems management across compute, storage, and networking resources in the Flex System infrastructure. It is designed to help you get the most out of your IT installation while automating repetitive tasks. It can significantly reduce the number of manual navigational steps for typical management tasks. From simplified system setup procedures with wizards and built-in expertise to consolidated monitoring for all of your physical and virtual resources, FSM provides core management functionality along with automation so that you can focus your efforts on business innovation.

The Flex System Manager compute node provides configuration and management support locally or remotely, for up to 16 chassis, their devices, and the compute nodes.

The following figure shows the Flex System Manager.



Figure 9. Flex System Manager

The Flex System Manager drives efficiency and cost savings in the data center. It provides a pre-integrated and virtualized management environment across servers, storage, and networking that is easily managed from a single pane of glass. It offers the following key features:

- Single focus point for seamless multichassis management provides an instant resource-oriented view of the chassis and chassis resources for compute nodes.
- Design reduces the number of interfaces, steps, and clicks that it takes to manage IT resources.
- Capability to intelligently manage and deploy workloads based on resource availability and predefined policies.
- Management of events and alerts to increase system availability and reduce downtime.
- It helps to drive energy savings and reduce operational costs.

The Flex System Manager offers these main features:

- Monitoring and problem determination:
 - A real-time multichassis view of hardware components with overlays for additional information
 - Automatic detection of issues in your environment through an event setup that triggers alerts and actions
 - Identification of changes that might affect availability
 - Server resource utilization by virtual machine or across a rack of systems
- Hardware management:
 - Automated discovery of physical and virtual servers and interconnections, applications, and supported third-party networking
 - Inventory of hardware components
 - Chassis and hardware component views, including hardware properties, component names and part numbers, firmware levels, energy usage and power consumption, utilization rates, and temperature readings
- Network management:
 - Management of network switches
 - Discovery, inventory, and status monitoring of switches
 - Graphical network topology views
 - Support for KVM, pHyp, and VMware virtual switches, and physical switches
 - Virtual LAN (VLAN) configuration of switches
 - Integration with server management
 - Per-virtual machine network usage and performance statistics provided to VMControl
 - Logical views of servers and network devices grouped by subnet and VLAN

- Storage management:

Storage can also be consolidated onto the FSM platform, or users can virtualize the integrated enterprise storage with other multivendor storage systems within the existing environment, and then control all virtualized storage resources from the Flex System platform.

- Virtualization and workload management:
 - Consolidate physical infrastructure or applications onto fewer servers
 - Deploy new or existing workload onto Linux, Microsoft Windows, or competitive platforms
 - Upgrade current rack or blade infrastructure to an integrated environment
 - Optimize the performance of applications by migrating them to Flex System Manager
 - Move into advanced stages of virtualization or a private cloud deployment
 - Deploy new workloads rapidly
 - Optimize or upgrade high-value, noninfrastructure workloads, such as enterprise resource planning (ERP), customer relationship management (CRM), data warehousing, or analytics

- Additional features:
 - Resource-oriented chassis map provides instant graphical view of chassis resources, including nodes and I/O modules
 - Remote console, including remote keyboard, video, mouse (KVM) session, remotely mounted optical and USB media, and remote power control
 - Hardware detection and inventory creation
 - Firmware compliance and updates
 - Automatic detection of hardware failures and alerting. Capability to take corrective action, including automatically notifying Lenovo of problems to escalate problem determination
 - Health status, such as CPU utilization, on all hardware devices from a single chassis view
 - Administrative capabilities, such as setting up users within profile groups, assigning security levels, and security governance

Power supplies

A maximum of six power supplies can be installed in the Enterprise Chassis. Either 2500W or 2100W power supplies are available. Two 2500W power supplies are standard in model 8721-A1x. Two 2100W power supplies are standard in model LRx. Additional power supplies are orderable per the following table. All installed power supplies must be the same.

Table 7. Ordering part number and feature code

Description	Part number	System x Feature code	Power Systems Feature code (7893-92X)
Flex System Enterprise Chassis 2500W Power Module (1 power supply)	43W9049	A0UC	3590
Flex System Enterprise Chassis 2100W Power Module (1 power supply)	47C7633	A3JH	None
Flex System Enterprise Chassis HVDC 2500W Power Module	00AM765	EPA9	EPA9
Flex System Enterprise Chassis -48V DC 2500W Power Module	00FJ635	A5VC	None

Each AC power supply part number ships with one 2.0m (6.5 ft) 16A/100-250V, C19 to IEC 320-C20 rack power cable, feature 6292. Other supported AC line cords including three-way split line cords, are listed in the following table.

Table 8. Supported AC line cords

Description	Part number	System x Feature code
4.3m, 16A/208V, C19 to NEMA L6-20P (US) Line Cord	40K9772	6275
2.5m, 16A/100-240V, C19 to IEC 320-C20 Rack Power Cable	39Y7916	6252
2m, 16A/100-250V, C19 to IEC 320-C20 Rack Power Cable	None	6292
4.3m, US/CAN, NEMA L15-30P - (3P+Gnd) to 3X IEC 320 C19	00D7192	A2Y3
4.3m, EMEA/AP, IEC 309 32A (3P+N+Gnd) to 3X IEC 320 C19	00D7193	A2Y4
4.3m, A/NZ, (PDL/Clipsal) 32A (3P+N+Gnd) to 3X IEC 320 C19	00D7194	A2Y5

The 2500W HVDC power supply operates at either 240V DC or 380V DC (192V - 400V input range). The power supply has a Rong-Feng RF-203P connector. The power supply includes a 2.5m DC power cord (SBB part number 00Y9100).

The -48V DC power supply operates at -48V to -60 V dc (nominal) and has a Molex 1060 Power Connector. The power supply ships with one 2.0m DC power cord (FRU part number 69Y1652)

The following figure shows the power supply and highlights the light-emitting diodes (LEDs).



Figure 10. Power supply option

Both power supplies are 80 PLUS Platinum-certified. The 2500W modules are 2500 Watts output rated at 200VAC to 208VAC (nominal), and 2750W at 220VAC to 240VAC (nominal). The power supply has an oversubscription rating of up to 3538 Watts output at 200VAC. The power supply operating range is 200-240 VAC. The power supplies also contain two dual independently powered 40mm cooling fan modules that are powered not from the power supply itself, but from the chassis midplane. The fan modules are variable speed and controlled by the chassis fan logic.

The 2100W power supplies are 2100 Watts output power rated at 200-240VAC. Similar to the 2500W unit, this power supply also supports oversubscription, the 2100W unit can run up to 2895W for short duration. The 2100W supplies have two independently powered dual 40mm cooling fans, that pick up power from the midplane included within the power supply assembly.

80 PLUS is a performance specification for power supplies used within servers and computers. To meet the 80 PLUS standard, the power supply must have an efficiency of 80% or greater, at 20%, 50%, and 100% of rated load with PF of .09 or greater. The standard has several grades, such as Bronze, Silver, Gold, and Platinum. More information on 80 PLUS is available at <http://www.80PLUS.org>.

The chassis allows configurations of power supplies to give N+N or N+1 redundancy. A chassis can operate on only three power-supply units (PSUs) with no redundancy, but N+1 or N+N is advised. Three power supplies (or six with N+N redundancy) allow for a balanced three-phase configuration.

All power supply modules are combined into a single power domain within the chassis, which distributes power to each of the compute nodes, I/O modules, and ancillary components through the Enterprise Chassis midplane. The midplane is a highly reliable design with no active components. Each power supply is designed to provide fault isolation and is hot swappable.

In the case of the 2500W supplies, power monitoring of both the DC and AC signals allows the Chassis Management Module to accurately monitor the power supplies. Power monitoring with the 2100W power supplies is currently not supported.

The following figure shows the compute node bay numbering (left) and power supply bay numbering (right).

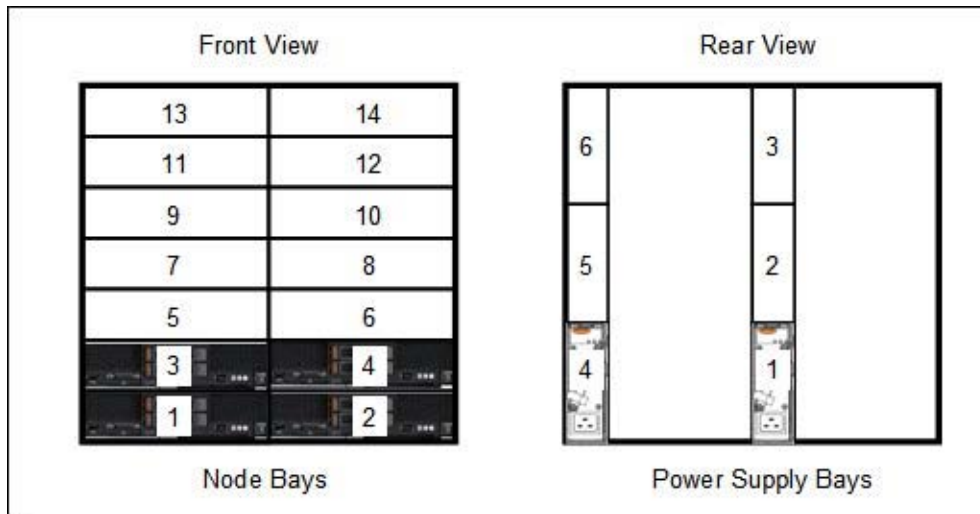


Figure 11. Power supply bay numbering

The following table shows the number of compute nodes that are supported in the chassis based on the power supplies used (2100W or 2500W), the number of power supplies installed, and the power redundancy policy enabled (N+1 or N+N).

In this table, the colors of the cells have the following meaning:

- Green cell: Supported with no restrictions as to the number of compute nodes that can be installed
- Yellow cell: Supported but with restrictions on the number of compute nodes that can be installed.

Note: These tables assume the chassis is filled with identical compute nodes. For more complex configurations, use the Power Configurator:

<http://ibm.com/systems/bladecenter/resources/powerconfig.html>

Table 9. Specific number of compute nodes supported based on installed power supplies (Part 1)

CPU TDP	2100W power supplies				2500W power supplies			
	N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total	N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total
Flex System x220								
50 W	14	14	14	14	14	14	14	14
60 W	14	14	14	14	14	14	14	14
70 W	14	14	14	14	14	14	14	14
80 W	14	14	14	14	14	14	14	14
95 W	14	14	14	14	14	14	14	14
Flex System x222								
50 W	14	14	13	14	14	14	14	14
60 W	14	14	12	12	14	14	14	14
70 W	14	14	11	12	14	14	14	14
80 W	14	14	10	11	14	14	13	14
95 W	14	13	9	10	14	14	12	13
Flex System x240								
60 W	14	14	11	11	14	14	14	14
70 W	14	14	11	11	14	14	14	14
80 W	14	14	11	11	14	14	14	14
95 W	14	14	11	11	14	14	14	14
115 W	14	14	11	11	14	14	14	14
130 W	14	14	11	11	14	14	13	14
135 W	14	14	11	11	14	14	13	14
Flex System x240 M5								
52 W	14	14	11	11	14	14	14	14
55 W	14	14	11	11	14	14	14	14
65 W	14	14	11	11	14	14	14	14
75 W	14	14	11	11	14	14	14	14
85 W	14	14	11	11	14	14	14	14
90 W	14	14	11	11	14	14	14	14
105 W	14	14	11	11	14	14	13	14
120 W	14	14	11	11	14	14	13	14
135 W	14	13	11	11	14	14	12	13
145 W	14	13	11	11	14	14	12	13

Table 9. Specific number of compute nodes supported based on installed power supplies (Part 2)

CPU TDP	2100W power supplies				2500W power supplies			
	N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total	N+1, N=5 6 total	N+1, N=4 5 total	N+1, N=3 4 total	N+N, N=3 6 total
Flex System x440								
95 W	7	7	6	6	7	7	7	7
115 W	7	7	5	5	7	7	7	7
130 W	7	7	5	5	7	7	6	7
Flex System x880 X6, x480 X6 and x280 X6								
105 W	7	7	5	5	7	7	7	7
130 W	7	7	5	5	7	7	7	7
155 W	7	7	5	5	7	7	7	7
Flex System V7000 Storage Node								
N/A	3	3	3	3	3	3	3	3
Flex System p24L								
All	14	12	9	10	14	14	12	13
Flex System p260								
All	14	12	9	10	14	14	12	13
Flex System p270								
All	14	12	9	9	14	14	12	12
Flex System p460								
All	7	6	4	5	7	7	6	6

Fan modules

The Enterprise Chassis supports up to a total of ten hot-swap fan modules: two 40 mm (1.57 in) fan modules and eight 80 mm (3.14 in) fan modules.

The two 40 mm fan modules distribute airflow to the I/O modules and chassis management modules. Both of these fan modules ship with the chassis. The 80 mm fan modules distribute airflow to the compute nodes through the chassis from front to rear. Each 80 mm fan module actually contains two 80 mm fan modules, back to back at each end of the module, which are counter-rotating. The following figure shows the 80 mm fan module.




Figure 12. 80 mm fan module

Four 80 mm fan modules are installed standard in chassis model 8721-A1x and CTO orders. The maximum number of 80 mm fan modules that can be installed is eight. Ordering information is shown in the following table. When the modules are ordered as a part number, they are supplied as a pair. Feature codes are for one fan.

Table 10. Fan module ordering information

Description	Part number	System x Feature code	Power Systems Feature code
Flex System Enterprise Chassis 80mm Fan Module	43W9078 (2 fan modules)	A0UA (1 fan module)	7805 (1 fan module)

The 80 mm fan modules are populated, depending on the nodes installed. To support the base configuration and up to four nodes, chassis model 8721-A1x ships with four 80 mm fan modules and two 40 mm fan modules preinstalled. There are two cooling zones for the nodes: a left zone and a right zone. Fan modules must be installed in pairs as shown in the following figure. If there are insufficient fan modules for the number of nodes installed, the compute nodes might be throttled to balance heat generation and cooling capacity.

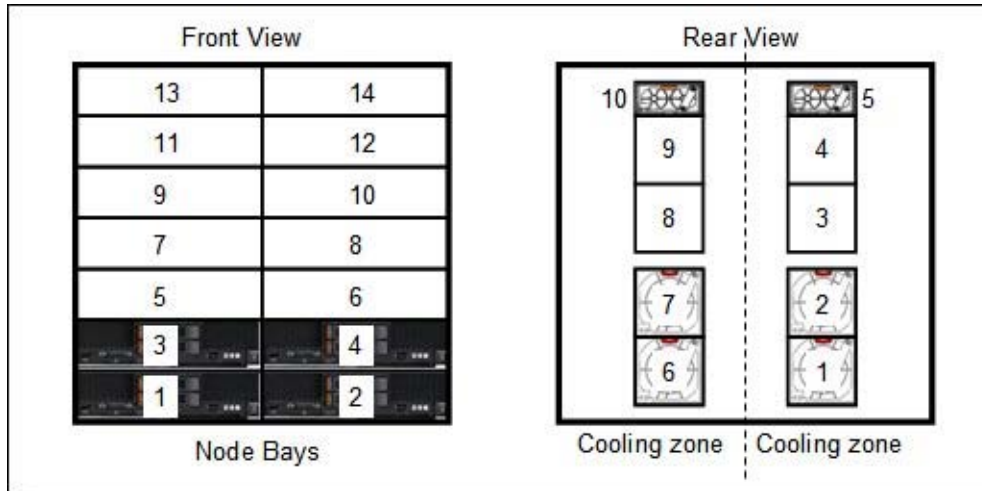


Figure 13. Fan module locations and cooling zones

The 40 mm fan modules are always required. Additional 80 mm fan modules are required as listed in the following table.

Table 11. 80 mm fan module requirements

Description	80 mm fan module requirements
Up to four half-wide compute nodes (node bays 1 - 4)	4 fan modules (fan bays 1, 2, and 6, 7)
Up to eight half-wide compute nodes (node bays 1 - 8)	6 fan modules (fan bays 1, 2, 3 and 6, 7, 8)
All 14 compute node bays (node bays 1 - 14)	8 fan modules (fan bays 1, 2, 3, 4 and 6, 7, 8, 9)

Physical specifications

Dimensions:

- Height: 440 mm (17.3 inches)
- Width: 447 mm (17.6 in)
- Depth, measured from front bezel to rear of chassis: 800 mm (31.5 in)
- Depth, measured from node latch handle to the power supply handle: 840 mm (33.1 in)

Weight:

- Minimum configuration: 96.62 kg (213 lb)
- Maximum configuration: 220.45 kg (486 lb)

Shipping dimensions (approximate):

- Height: 81 cm (32 in)
- Length: 102 cm (40 in)
- Width: 71 cm (24 in)

Supported environment

The Flex System Enterprise Chassis complies with ASHRAE Class A3 specifications.

The following environment is the supported operating environment.

Temperature:

- 0 - 914 m (0 - 3,000 ft): 5 - 40 °C (41 - 104 °F)
- 914 m - 3048 m (3,000 - 10,000 ft): The maximum ambient temperature drops 1 °C for every additional 178 m (584 ft) increase in altitude until the maximum temperature is 28 °C at 3,048 m (10,000 ft)

Relative humidity: 8% - 85%

Maximum altitude: 3,048 m (10,000 ft)

Electrical power:

- 200 - 240 V ac (nominal), 50 or 60 Hz
- Minimum configuration: 0.51 kVA (two power supplies)
- Maximum configuration: 13 kVA (six power supplies)

Power consumption: 12,900 watts maximum

Thermal output, full configuration - 12,900 watts (43,900 Btu/hr)

Acoustical noise emissions for Flex Chassis:

- 7.5 bels operating
- 7.5 bels idling

The noise emission level stated is the declared (upper limit) sound power level, in bels, for a random sample of machines. All measurements are made in accordance with ISO 7779 and reported in conformance with ISO 9296.

Warranty options

The Flex System Enterprise Chassis has a three-year on-site warranty with 9x5 next-business-day terms. Lenovo offers the warranty service upgrades through ServicePac®, discussed in this section. The ServicePac is a series of prepackaged warranty maintenance upgrades and post-warranty maintenance agreements with a well-defined scope of services, including service hours, response time, term of service, and service agreement terms and conditions.

ServicePac offerings are country-specific. That is, each country might have its own service types, service levels, response times, and terms and conditions. Not all covered types of ServicePac might be available in a particular country. For more information about ServicePac offerings available in your country, see the ServicePac Product Selector at <https://www-304.ibm.com/sales/gss/download/spst/servicepac>.

The following table explains warranty service definitions in more detail.

Table 12. Warranty service definitions

Term	Description
On-site repair	A service technician will come to the server's location for equipment repair.
24x7x2 hour	A service technician is scheduled to arrive at your client's location within two hours after remote problem determination is completed. We provide 24-hour service, every day, including holidays.
24x7x4 hour	A service technician is scheduled to arrive at your client's location within four hours after remote problem determination is completed. We provide 24-hour service, every day, including holidays.
9x5x4 hour	A service technician is scheduled to arrive at your client's location within four business hours after remote problem determination is completed. We provide service from 8:00 a.m. to 5:00 p.m. in the client's local time zone, Monday through Friday, excluding holidays. If after 1:00 p.m., it is determined that on-site service is required, the client can expect the service technician to arrive the morning of the following business day. For noncritical service requests, a service technician will arrive by the end of the following business day.
9x5 next business day	A service technician is scheduled to arrive at your client's location on the business day after we receive your call, following remote problem determination. We provide service from 8:00 a.m. to 5:00 p.m. in the client's local time zone, Monday through Friday, excluding holidays.

In general, the following types are the types of ServicePacs:

- Warranty and maintenance service upgrades:
 - One, two, three, four, or five years of 9x5 or 24x7 service coverage
 - On-site repair from the next business day to four or two hours
 - One or two years of warranty extension
- Remote technical support services:
 - One or three years with 24x7 coverage (severity 1) or 9 - 5 next business day for all severities
 - Installation and start-up support for System x servers
 - Remote technical support for System x servers
 - Software support - Support Line:
 - Microsoft or Linux software
 - VMware
 - IBM Systems Director

Regulatory compliance

The server conforms to the following standards:

- ASHRAE Class A3
- FCC - Verified to comply with Part 15 of the FCC Rules Class A
- Canada ICES-004, issue 3 Class A
- UL/IEC 60950-1
- CSA C22.2 No. 60950-1
- NOM-019
- Argentina IEC 60950-1
- Japan VCCI, Class A
- IEC 60950-1 (CB Certificate and CB Test Report)
- China CCC (GB4943); (GB9254, Class A); (GB17625.1)
- Taiwan BSMI CNS13438, Class A; CNS14336
- Australia/New Zealand AS/NZS CISPR 22, Class A
- Korea KN22, Class A, KN24
- Russia/GOST ME01, IEC 60950-1, GOST R 51318.22, GOST R 51318.249, GOST R 51317.3.2, GOST R 51317.3.3
- CE Mark (EN55022 Class A, EN60950-1, EN55024, EN61000-3-2, EN61000-3-3)
- CISPR 22, Class A
- TUV-GS (EN60950-1/IEC 60950-1, EK1-ITB2000)

External disk storage systems

The following System x options are available for attaching external storage systems to Enterprise Chassis:

Table 13. External disk storage systems

Part number	Description
2071CU2	IBM Storwize V3500 LFF Dual Control Enclosure
2071CU3	IBM Storwize V3500 SFF Dual Control Enclosure
6099L2C	IBM Storwize V3700 3.5-inch Storage Controller Unit
6099S2C	IBM Storwize V3700 2.5-inch Storage Controller Unit
6099T2C	IBM Storwize V3700 2.5-inch DC Storage Controller Unit
6194L2C	IBM Storwize V5000 LFF Control Enclosure
6194LEU	IBM Storwize V5000 LFF Expansion Enclosure
6194S2C	IBM Storwize V5000 SFF Control Enclosure
6194SEU	IBM Storwize V5000 SFF Expansion Enclosure
6195SC5	IBM Storwize V7000 2.5-inch Storage Controller Unit
6195LEF	IBM Storwize V7000 3.5-inch Storage Expansion Unit
6195SEF	IBM Storwize V7000 2.5-inch Storage Expansion Unit

Top-of-rack Ethernet switches

For enterprise-class installations with multiple Flex System Enterprise Chassis configurations, a top-of-rack Ethernet switch from Lenovo provides the necessary level of networking between racks of systems and the rest of your production network. The following table lists the available top-of-rack switches.

Table 14. Top-of-rack switches

Part number	Description
1 Gb top-of-rack switches	
7309BAX	RackSwitch G7028 (24 ports)
7309CAX	RackSwitch G7028 (48 ports)
0446013	RackSwitch G8000R
7309CFC	RackSwitch G8000F
7309G52	RackSwitch G8052R
730952F	RackSwitch G8052F
10 Gb top-of-rack switches	
7309DRX	RackSwitch G8264CS (Rear to Front)
7309DFX	RackSwitch G8264CS (Front to Rear)
7309BR6	RackSwitch G8124ER
7309BF7	RackSwitch G8124EF
7309G64	RackSwitch G8264R
730964F	RackSwitch G8264F
7309CR9	RackSwitch G8264TR
7309CF9	RackSwitch G8264TF
40 Gb top-of-rack switches	
8036BRX	RackSwitch G8332 (Rear to Front)
8036BFX	RackSwitch G8332 (Front to Rear)
8036ARX	RackSwitch G8316R
8036AFX	RackSwitch G8316F

Power distribution units

Power planning for an Flex System Enterprise Chassis is essential.

The Enterprise Chassis has a maximum of six power supplies installed. So, careful consideration must be given to providing the best power-optimized source. Both N+N and N+1 configurations are supported for maximum flexibility in power redundancy. Each AC power supply in the chassis has a 16A C20 three-pin socket and can be fed by a C19 power cable, from a suitable supply.

The chassis has the ability to accommodate a maximum of six power supplies, so it is possible to balance a three-phase power input into a single chassis or a group of chassis. The chassis power system is designed for efficiency using data center power consisting of three-phase, 60A Delta 200 VAC (North America) or three-phase 32A wye 380-415 VAC (international). The chassis can also be fed from single-phase 200 - 240VAC supplies, if required.

For further details, see the *Flex System Enterprise Chassis Power Requirements Guide*, available from: <http://ibm.com/systems/bladecenter/resources/powerconfig.html>

Table 15. Power distribution units

Part number	Description
Switched and monitored AC PDUs	
46M4002	1U 9 C19/3 C13 Active Energy Manager DPI® PDU
46M4003	1U 9 C19/3 C13 Active Energy Manager 60A 3 Phase PDU
46M4167	1U 9 C19/3 C13 Switched and Monitored 30A 3 Phase PDU
46M4134	0U 12 C19/12 C13 Switched and Monitored 50A 3 Phase PDU
46M4140	0U 12 C19/12 C13 50A 3 Phase PDU
Enterprise AC PDUs	
71762MX	Ultra Density Enterprise PDU C19 PDU+ (WW)
71762NX	Ultra Density Enterprise PDU C19 PDU (WW)
71763MU	Ultra Density Enterprise PDU C19 3 phase 60A PDU+ (NA)
71763NU	Ultra Density Enterprise PDU C19 3 phase 60A PDU (NA)
39Y8923	DPI 60A Three Phase C19 Enterprise PDU with IEC309 3P+G (208 V) fixed line cord
39Y8948	DPI Single Phase C19 Enterprise PDU without line cord
Front-end AC PDUs	
39Y8934	DPI 32amp/250V Front-end PDU with IEC 309 2P+Gnd connector
39Y8938	30amp/125V Front-end PDU with NEMA L5-30P connector
39Y8939	30amp/250V Front-end PDU with NEMA L6-30P connector
39Y8940	60amp/250V Front-end PDU with IEC 309 60A 2P+N+Gnd connector
DC PDUs	
44T0966	1U Higher Voltage DC PDU (240V/380V)

Uninterruptible power supply units

The Flex System Enterprise Chassis supports attachments to the uninterruptible power supply units listed in the following table.

Table 16. Uninterruptible power supply units

Part number	Description
55946KX	RT6kVA 3U Rack or Tower UPS (200-240VAC)
55948KX	RT8kVA 6U Rack or Tower UPS (200-240VAC)
55949KX	RT11kVA 6U Rack or Tower UPS (200-240VAC)
55948PX	RT8kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)
55949PX	RT11kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)
21303RX	Uninterruptible Power Supply 7500XHV
21304RX	Uninterruptible Power Supply 10000XHV
53956AX	6000VA LCD 4U Rack Uninterruptible Power Supply (200 V/208 V)
53956KX	6000VA LCD 4U Rack Uninterruptible Power Supply (230 V)
53959KX	11000VA LCD 5U Rack Uninterruptible Power Supply (230V)

For more information, see the related Lenovo Press Product Guides in the Power infrastructure category:
<http://www.redbooks.ibm.com/portals/systemx?Open&page=pg&cat=power>

Rack cabinets

The Flex System Enterprise Chassis is supported in the rack cabinets as listed in the following table. The suggested rack for Flex System is the 42U 1100 mm Enterprise V2 Deep Dynamic Rack, 93634PX.

Table 17. Rack cabinets

Part number	Feature code	Rack cabinet	Supported
93634PX	A1RC	42U 1100 mm Enterprise V2 Deep Dynamic Rack	Recommended
93634EX	A1RD	42U 1100 mm Dynamic Enterprise V2 Expansion Rack	Recommended
93634CX	A3GR	PureFlex System 42U Rack	Recommended
93634DX	A3GS	PureFlex System 42U Expansion Rack	Recommended
93634AX	A31F	PureFlex System 42U Rack	Recommended
93634BX	A31G	PureFlex System 42U Expansion Rack	Recommended
201886X	2731	11U Office Enablement Kit	Yes*
93072PX	6690	S2 25U Static Standard Rack	Yes
93072RX	1042	S2 25U Dynamic Standard Rack	Yes
93074RX	1043	S2 42U Standard Rack	Yes
99564RX	5629	S2 42U Dynamic Standard Rack	Yes
99564XX	5631	S2 42U Dynamic Standard Expansion Rack	Yes
93084PX	5621	42U Enterprise Rack	Yes
93084EX	5622	42U Enterprise Expansion Rack	Yes
93604PX	7649	42U 1200 mm Deep Dynamic Rack	Yes
93604EX	7650	42U 1200 mm Deep Dynamic Expansion Rack	Yes
93614PX	7651	42U 1200 mm Deep Static Rack	Yes
93614EX	7652	42U 1200 mm Deep Static Expansion Rack	Yes
93624PX	7653	47U 1200 mm Deep Static Rack	Yes
93624EX	7654	47U 1200 mm Deep Static Expansion Rack	Yes
14102RX	1047	eServer™ Cluster 25U Rack	Yes
14104RX	1048	Linux Cluster 42U Rack	Yes
9306-900	None	Netfinity® Rack	No
9306-910	None	Netfinity Rack	No
9306-42P	None	Netfinity Enterprise Rack	No
9306-42X	None	Netfinity Enterprise Rack Expansion Cabinet	No
9306-200	None	Netfinity NetBAY 22	No

* This Office Enablement kit is specifically designed for the BladeCenter S Chassis. The Flex System Enterprise Chassis can be installed in this kit with 1U of space remaining, however the noise generated by the chassis may not be acceptable for office use.

Rack options

The server supports the rack console switches and monitor kits listed in the following table.

Table 18. Rack options

Part number	Feature code	Description
Monitor kits and keyboard trays		
17238BX	1723HC1 fc A3EK	1U 18.5" Standard Console
17238EX	1723HC1 fc A3EL	1U 18.5" Enhanced Media Console
172317X	1723HC1 fc 0051	1U 17in Flat Panel Console Kit
172319X	1723HC1 fc 0052	1U 19in Flat Panel Console Kit
Console switches		
3858D3X	3858HC1 fc A4X1	Avocent Universal Management Gateway 6000 for IBM
1754D2X	1754HC2 fc 6695	Global 4x2x32 Console Manager (GCM32)
1754D1X	1754HC1 fc 6694	Global 2x2x16 Console Manager (GCM16)
1754A2X	1754HC4 fc 0726	Local 2x16 Console Manager (LCM16)
1754A1X	1754HC3 fc 0725	Local 1x8 Console Manager (LCM8)
Console cables		
00AK142	A4X4	UM KVM Module VGA+SD Dual RJ45
43V6147	3757	Single Cable USB Conversion Option (UCO)
39M2895	3756	USB Conversion Option (4 Pack UCO)
39M2897	3754	Long KVM Conversion Option (4 Pack Long KCO)
46M5383	5341	Virtual Media Conversion Option Gen2 (VCO2)
46M5382	5340	Serial Conversion Option (SCO)

For more information, see the related Lenovo Press Product Guides in the Power infrastructure category:
<http://www.redbooks.ibm.com/portals/systemx?Open&page=pg&cat=rack>

Related publications and links

For more information, see the following resources:

- Flex System product page
<http://ibm.com/systems/flex>
- US Product Announcement for the Flex System Enterprise Chassis
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS112-053>
- Lenovo Press Product Guides for Flex System servers and options
<http://www.redbooks.ibm.com/portals/FlexSystem?Open&page=pgbycat>
- Flex System Information Center
<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>
- *Flex System Enterprise Chassis Installation and Service Guide*
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8721.doc/printable_doc.html
- ServerProven hardware compatibility page for Flex System
<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>
- *Flex System Interoperability Guide*
<http://www.redbooks.ibm.com/fsig>
- *Flex System Enterprise Chassis Power Requirements Guide*
<http://ibm.com/systems/bladecenter/resources/powerconfig.html>
- Lenovo Press publication *Flex System Products and Technology by Lenovo*, SG24-8255
<http://www.redbooks.ibm.com/abstracts/sg248255.html>
- IBM Configurator for e-business (e-config)
<http://ibm.com/services/econfig/>
- System x and Cluster Solutions configurator (x-config)
<http://ibm.com/products/hardware/configurator/americas/bhui/asit/>
- System x Configuration and Options Guide:
<http://ibm.com/systems/xbc/cog/>
- ServerProven for Flex System
<http://ibm.com/systems/info/x86servers/serverproven/compat/us/flexsystems.html>
- IBM System Storage® Interoperation Center
<http://www.ibm.com/systems/support/storage/ssic>

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing*

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2012-2015. All rights reserved.

This document was created or updated on January 29, 2015.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbooks@us.ibm.com

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips0863.html> .

Trademarks

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

The following terms are trademarks of Lenovo in the United States, other countries, or both:

BladeCenter®
Flex System™
Lenovo®
Netfinity®
RackSwitch™
Lenovo (logo)®
ServerProven®
System x®
TopSeller™

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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