



## Product Guide

April 2011

# IBM System x3690 X5

## Architecture Overview

CONTENTS	
Architecture Overview	1
Key Highlights	2
Key Features	3
Key Options	13
x3690 X5 Images	14
x3690 X5 Specs	16
The Bottom Line	19
Server Comparison	21
For More Information	22
Legal Information	22

Outstanding performance, superior mainframe-like reliability, and fault-tolerant memory characteristics

IBM® has been designing and implementing servers under the X-Architecture® name since 2001. eX5 technology represents the fifth generation of enterprise servers based on the same design principle IBM began with in 1997: to offer systems that are expandable, offer “big iron” reliability, availability, and serviceability (RAS) features, with extremely competitive price/performance on an Intel® Xeon® processor-based system.

The eX5 technology is primarily designed around three major workloads: database servers, server consolidation using virtualization services, and Enterprise Resource Planning (application and database) servers.

If you're using industry-standard Xeon processor-based servers for running business critical applications, the systems that run these applications need the type of technology designed into IBM's eX5 technology systems. eX5 technology represents a \$100M+ investment in designing a flagship offering that can harness the power of the latest 64-bit x86 (x64) Xeon processors. The eX5 family includes a **2U 2-socket** rack-optimized chassis (**x3690 X5**). In addition, for those workloads that require the maximum available memory, the x3690 X5 support a 1U **MAX5** memory expansion unit, which adds an additional **32 DIMM slots** (up to **64 DIMMs / 2TB** total). You can save money on software licensing by using the memory capacity to virtualize the server into many VMs, rather than using multiple servers. Huge amounts of memory also enable larger VMs and larger databases.

The x3690 X5 offers database optimized models with eXFlash solid-state drive technology, pretested and ready to deploy for database workloads.

The **x3690 X5 Workload Optimized Solution for SAP In-Memory Appliance**, SAP HANA models offer a preloaded and optimized appliance based on the two-socket x3690 X5 server platform. These models include two processors, 256GB of memory and choice of all eXFlash SSs or a combination of solid-state and hard disk drives. They are designed for use in small-to-midsized SAP HANA configurations. These models also include the following software: Novell SLES for SAP Applications operating system with three-year priority support and the IBM General Parallel File System (GPFS) with 3 year support. **Note:** SAP HANA software is included, but sold separately by SAP.

The x3690 X5 offers models **integrated** with a **hypervisor preloaded** on a **USB 2.0 flash key**. It operates in a diskless configuration, offers a smaller memory footprint, extremely high performance, and stronger security, making getting a system up and running in a virtualized environment faster and easier than ever before.

IBM X-Architecture pioneered XpandOnDemand™ (“pay as you grow”) scalability, which allows chassis to be simply cabled together to form larger scale-up systems. This capability allows IBM to sell a large SMP (symmetric multiprocessing) system at entry price points. With XpandOnDemand, you can start small and later expand as your needs change, without requiring you to buy more than you need up-front or throw away parts later as you expand.

IBM's eX5 technology-based systems are the ideal solution for scale-up database-serving applications on Microsoft® Windows® with Microsoft SQL Server® or IBM DB2®, as well as Linux® with Oracle or DB2. Database hosting demands ultimate server reliability features, and once installed, they tend to grow and grow, requiring ever greater levels of availability. eX5 servers provide exactly that degree to availability.

Another strong application area for the eX5-based systems is enterprise server consolidation activities workloads, including SAP and Oracle. eX5 systems can offer considerable savings over UNIX® deployments, using our certified solution stacks on either Windows or Linux.. Larger servers need more processor, memory and I/O resources, which make maximum use of any applicable virtual machine software licensing fees and deliver superior system utilization levels. The name of the game in consolidation activities is to deploy the fewest new servers possible and help IT staff manage more images with the same or fewer overall people.

The eX5 servers are designed to protect your data with high performance, high reliability, and high availability. They support the latest **ten-, eight-, six-, and four-core** Intel® Xeon® processor **E7** family and **7500/6500** series processors, up to **1066MHz** memory access, and up

to **30MB** of **L3** cache to help provide you with the computing power you need to match your business needs and growth. The new Xeon **E7-8800**, **E7-4800**, and **E7-2800** series processors offer up to **40%** greater performance than the previous-generation **7500** and **6500** series processors.

In addition, the x3690 X5 uses industry-standard **DDR3** memory with **Chipkill™ ECC** (Error Checking and Correcting) protection—for high performance and reliability. For even higher levels of availability, the eX5 servers also offer **Memory ProteXion™**, **memory scrubbing**, **memory rank sparing**, and selectable **memory mirroring**.

A **dual-port** integrated high-speed **Gigabit Ethernet** controller with **TOE** (TCP Offload Engine) acceleration is standard in all models, as are **five** high-performance **PCIe** adapter slots. A **10 Gigabit Ethernet** virtual fabric adapter is standard in selected models. The 10GbE solution provides the benefits and flexibility of I/O convergence in a single end-to-end solution.

The x3690 X5 offer industry-leading scalability, including up to **1TB** of memory in the 2U chassis (upgradeable to **2TB** using a 1U MAX5 memory expansion unit), and a choice of up to **16 2.5-inch** internal **Serial-Attach SCSI** (SAS) or **Serial ATA II** (SATA II) **hot-swap** hard disk drives, with an internal storage capacity of up to **9.6TB (SAS)** or **8TB (SATA)**, or up to **24 1.8-inch solid-state drives (SSDs)** with a capacity of **4.8TB**. Or, you can combine the two, having **8 HDDs and 16 SSDs** in the same chassis. The Standard **ServeRAID-M1015** SAS/SATA controller provides RAID-0/1/10 support (upgradeable to 5/50). Optional **IBM ServeRAID** SAS/SATA/SSD RAID and SSD non-RAID controllers add up to **512MB** of battery-backed cache to the onboard controller to provide **RAID-0/1/10/10/5/50/6/60**. The **2U** size of the chassis helps you maximize your rack investments. Up to **21** of these chassis can be installed in a single 42U rack, for a total of up to **42** processors (**420** cores), **105** PCIe slots, and **336** HDDs or **504** SSDs, offering an ideal balance of performance, storage and I/O slots per rack.

Standard in the x3690 X5 is the **Integrated Management Module (IMM)** that enables the user to manage and control the server easily—both locally and remotely. The IMM offers a high level of manageability that is designed to keep costs down and the system up—even when network usage increases. The drop-down **light path diagnostics** panel enables quick servicing of the system if a problem develops. These advanced features help maximize network availability by increasing uptime, as do IBM OnForever™ features, such as **hot-swap/redundant HDDs, power and fans; Active Memory™**; **temperature-controlled fans**; **IPMI 2.0** support, including **highly secure remote power control** and **Serial over LAN**; as well as **text- and graphics-console redirect over LAN**.

Another improvement with the new generation of X-Architecture is the replacement of legacy BIOS with a new generation **Unified Extensible Firmware Interface (UEFI)**. UEFI provides a more intuitive user interface and understandable event logs and better management.

With the inclusion of unique IBM service and support features such as **IBM Systems Director**, **IBM Systems Director Active Energy Manager™**, and **IBM ServerGuide™**, the x3690 X5 is designed for maximum uptime.

If you need a balance of high-performance two-socket processing, enormous memory expansion, and large I/O capacity in a rack-dense environment, these are the ideal systems.

---

## Key Highlights

## eX5 Design Features

eX5 design offers numerous features to boost performance and reduce product and operating costs:

- Supports up to **two** Intel Xeon processors with **12MB** to **30MB** of **L3** cache (processor dependent), and Intel Virtualization Technology (VT), Hyper-Threading (HT) technology, and Turbo Boost technology.
- Supports **10-core**, **8-core**, or **6-core Xeon E7** family processors or **8-core**, **6-core** or **4-core** Xeon 7500 series processors.
- Supports up to **1TB** of **DDR3** memory in a single 2U chassis/node.
- The 7147 machine type supports low-voltage **1.35V** DIMMs.
- Support for a **MAX5** memory expansion unit, increasing memory capacity by **100%** for memory-intensive workloads. One chassis with a MAX5 supports up to **2TB** of memory in 3U<sup>1</sup>.
- Alternatively, with 64 DIMM slots, the server and MAX5 can be populated with inexpensive 2GB DIMMs (for 128GB), while gaining the performance boost from using **8 memory controllers** (4 in the server and four more in the MAX5 unit).
- Fast **PC3-10600 DDR3** ECC memory with **Chipkill** error correction, **Memory ProteXion™** (redundant bit steering), **memory scrubbing**, **memory rank sparing**, and selectable


---

<sup>1</sup> 7147 machine type; 7148 machine type is limited to 1TB (512GB in the chassis and 512GB in the MAX5)..

**memory mirroring** protection provides speed and high availability.

- **Five** high-speed **PCIe Gen 2** adapter slots offer investment protection by supporting high-performance adapters, such as 10Gb Ethernet, Fibre Channel, Converged Network Adapter (CNA), and InfiniBand™ adapters.
  - The flexibility to configure a **2U** or **3U** system (using 1 server and 1 MAX5 unit).
  - **2-Port 10Gb Ethernet** virtual fabric adapter and **2-port Gigabit Ethernet** controller standard in selected models.
  - **Integrated Management Module** (IMM) standard.
  - Optional **ServeRAID** controllers that support **RAID-0/1/10/5/50/6/60** and high-IOPS **SSD** drives.
  - Up to **9.6TB** of internal high-speed **hot-swap SAS** storage, or **8TB** of low-cost **hot-swap SATA** storage, using **16 2.5-inch HDDs**; up to **4.8TB** of **hot-swap solid-state storage**, using **24 SSDs**.
- 

## Key Features



## Multicore Xeon Processors

The x3690 X5 supports up to **two** high-performance Intel **Xeon E7** family or **7500/6500** series processors (depending on machine type), allowing you to upgrade as business needs require.

The system offers a choice of processor clock rates, cache sizes and power draw:

### Xeon E7 Family Processors

- **105W 10-core** Xeon processor model **E7-8867L** running at 2.13GHz, with *extremely low power draw per core (10.5W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **30MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology—processor available via CTO only
- **130W 8-core** Xeon processor model **E7-8837** running at 2.67GHz, with *reduced power draw per core (16.25W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **105W 6-core** Xeon processor models **E7-4807\*** running at 1.86GHz, with *moderate power draw per core (17.5W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, HT, and HT technology (no Turbo Boost)—processor available via CTO only
- **130W 10-core** Xeon processor models **E7-2870** running at 2.4GHz, with *low power draw per core (13W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **30MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **130W 10-core** Xeon processor models **E7-2860** or **E7-2850** running at 2.26 or 2.0GHz (respectively), with *low power draw per core (13W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **105W 8-core** Xeon processor model **E7-2830** running at 2.13GHz, with *low power draw per core (13.13W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology—processor available via CTO only
- **105W 8-core** Xeon processor models **E7-2820** running at 2.0GHz, with *low power draw per core (13.13W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **105W 6-core** Xeon processor models **E7-2803** running at 1.73GHz, with *moderate power draw per core (17.5W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, and HT technology (no Turbo Boost)

### Xeon 7500 and 6500 Series Processors

- **130W 8-core** Xeon processor model **X7560** running at 2.26GHz, with *low power draw per core (16.25W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
  - **130W 8-core** Xeon processor model **X7550** running at 2.0GHz, with *low power draw per core (16.25W)*, **6.4 GTps** QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
  - **130W 8-core** Xeon processor model **X6550** running at 2.0GHz, with *low power draw per core (16.25W)*, **6.4 GTps** QPI speed, **1066MHz** memory access, dual integrated memory
-

controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology

- **95W 8-core** Xeon processor model **L7555** running at 1.86GHz, with *extremely low power draw per core (11.88W)*, **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **95W 6-core** Xeon processor model **L7545** running at 1.86GHz, with *low power draw per core (15.83W)*, **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **130W 6-core** Xeon processor model **L7542** running at 2.66GHz, with *reduced power draw per core (21.67W)*, **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT and Turbo Boost technology (no HT)
- **105W 6-core** Xeon processor model **E7540** running at 2.0GHz, with *low power draw per core (17.5W)*, **6.4** GTps QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **105W 6-core** Xeon processor model **E6540** at 2.0GHz, with 64-bit extensions, *low power draw per core (17.7W)*, **5.86** GTps QPI speed, **1066MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology
- **105W 6-core** Xeon processor model **E7530** at 1.86GHz, with 64-bit extensions, *low power draw per core (17.5W)*, **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **12MB** of shared **L3** cache, and Intel VT, HT, and Turbo Boost technology
- **95W 4-core** Xeon processor model **E7520** at 1.86GHz, with 64-bit extensions, *reduced power draw per core (23.75W)*, **4.8** GTps QPI speed, **800MHz** memory access, dual integrated memory controllers, **12MB** of shared **L3** cache, and Intel VT and Hyper-Threading technology (no Turbo Boost)
- **105W 4-core** Xeon processor model **E6510\*** at 1.73GHz, with 64-bit extensions, *reduced power draw per core (26.25W)*, **4.8** GTps QPI speed, **800MHz** memory access, dual integrated memory controllers, **12MB** of shared **L3** cache, and Intel Hyper-Threading technology

\* This processor is limited to single-chassis memory, *without* MAX5.

With the Xeon E7 family, and 7500/6500 series processors, Intel has diverged from its traditional Symmetric Multiprocessing (SMP) architecture to a Non-Uniform Memory Access (NUMA) architecture. The processors are connected through three serial coherency links called QuickPath Interconnect (QPI). QPI is capable of **6.4 GTps** (gigatransfers per second). Put another way, the E7/7500/6500 processors offer **3X** the bandwidth of the previous 7400 series processors.

**Turbo Boost Technology** dynamically turns off unused processor cores and increases the clock speed of the cores in use, by up to three model frequencies. For example, with **7-10** cores active, a 2.4GHz E7-2870 10-core processor can run the cores at up to **2.53GHz**. With **5-6** cores active, it can run those cores at **2.67GHz**; with only **1-4** cores active, it can run those cores at **2.8GHz**. Similarly, with six cores active, a 2.26GHz X7560 8-core processor can run the cores at **2.53GHz**. With only 3-4 cores active, it can run those cores at **2.67GHz**. When the cores are needed again, they are dynamically turned back on and the processor frequency is adjusted accordingly.

Each processor includes two integrated memory controllers, to reduce memory bottlenecks and improve performance. Memory access is at up to **1066MHz** frequency, depending on the processor model and memory used.

In processors implementing **Hyper-Threading technology**, each core has two threads capable of running an independent process. Thus, an 8-core processor can run **16** threads concurrently.

Intel's **Virtualization Technology** (VT) integrates hardware-level virtualization hooks that allow operating system vendors to better utilize the hardware for virtualization workloads.

**Intelligent Power Capability** powers individual processor elements on and off as needed, to reduce power draw.

---

## DDR3 Registered Memory with Active Memory Protection

eX5 servers use registered double data rate III (DDR3) DIMMs and provides Active Memory features, including advanced **Chipkill** memory protection, for **up to 16X** better error correction than standard ECC memory, **IBM Memory ProteXion**, and optional **memory mirroring** or **memory rank sparing**. In addition to offering **triple the memory bandwidth** of registered DDR2 or fully-buffered memory, DDR3 memory also uses less energy. 1.8V DDR2 memory already offered up to **37%** lower energy use than fully buffered memory. Now, a generation later,


---

1.5V DDR3 memory is even more efficient, using **10-15% less energy** than DDR2 memory. In addition machine type 7147 supports **low-voltage 1.35V** DIMMs, which consume up to **19%** less energy than 1.5V DIMMs. (1.5V and 1.35V DIMMs can be mixed within a given memory controller, however all DIMMs in that controller will run at 1.5V.)

The **7147** machine type of x3690 X5 chassis supports up to **1TB** of memory in **32** DIMM sockets using **32GB** DIMMs. (The **7148** machine type supports up to **512GB** with **16TB** DIMMs.) It uses **PC3-10600** double data rate 3 (DDR3) memory (operating at **1066MHz**) for fast access, as well as PC3-8500 DDR3 and LV versions of both types. Competitive servers may require you to add processors to add memory. eX5 servers with MAX5 decouple memory and processors, so that you can upgrade memory without upgrading to a second processor.

The standard configuration supports up to **32** DIMMs. If that is not enough, you can attach a 1U **MAX5** memory expansion unit, which adds another **32 DIMM slots** (up to **1TB** in machine type 7147, for a total of **2TB** in only **3U**).

Redesign in the architecture of the 7500/6500 series and E7 processors bring radical changes in the way memory works in these servers. For example, the processors **integrate two memory controllers inside each processor**, resulting in **four** memory controllers in a two-socket system. (MAX5 adds another **four** memory controllers.) Each controller has two memory channels (four per processor), each with a memory buffer.



The memory architecture of the eX5 chipset provides up to **333%** more aggregate memory bandwidth than the previous generation eX4, for exceptional memory performance. By performing reads and writes simultaneously, it eliminates memory read-to-write blocking latency. In addition, it also offers innovative data reliability and security features to help improve data integrity, including enhanced CRC protection, data retry on error detect and buffer registers for improved fault isolation.

It is important to ensure that DIMMs with appropriate number of ranks are populated in each channel for optimal performance. Whenever possible, **it is recommended to use dual-rank DIMMs** in the system. Dual-rank DIMMs offer better interleaving and hence better performance than single-rank DIMMs.

Another important guideline is to populate equivalent ranks per channel. For instance, **mixing one single-rank DIMM and one dual-rank DIMM in a channel should be avoided**.

**Notes:** DIMMs must be installed in matching pairs. Also, each CPU requires at least 2 DIMMs. It is important to ensure that all memory channels in each processor are populated. The relative memory bandwidth decreases as the number of channels populated decreases. This is because



the bandwidth of all the memory channels is utilized to support the capability of the processor. So, as the channels are decreased, the burden to support the requisite bandwidth is increased on the remaining channels, causing them to become a bottleneck.

**For peak performance:**

- Populate both processors with equal amounts of memory to enable a balanced NUMA system.
- Populate both memory controllers within a processor with equal memory capacity.
- Populate an even number of ranks per channel.
- Use dual-rank DIMMs whenever appropriate.

**Power guidelines:**

- Fewer larger DIMMs (for example 8 x 4GB DIMMs vs. 16 x 2GB DIMMs will generally have lower power requirements
- x8 DIMMs (x8 data width of rank) will generally draw less power than equivalently sized x4 DIMMs

**Reliability guidelines:**

- Using fewer, larger DIMMs (for example 8 x 4 GB DIMMs vs. 16 x 2GB DIMMs is generally more reliable
- Memory controllers support IBM Chipkill memory protection technology with x4 DIMMs (x4 data width of rank) and x8 DIMMs

For increased availability beyond Chipkill error correction, the eX5 servers offer additional levels of IBM Active Memory protection:

**Memory ProteXion** technology provides *multichip* error protection and works in conjunction with **Chipkill** technology—which provides *multibit* protection per chip—and standard ECC protection, to provide multi-level memory correction.

**Memory scrubbing** is an automatic daily test of all system memory. It detects and reports memory errors that might be developing before they cause a server outage. Memory scrubbing and Memory ProteXion work together. When a bit error is detected, memory scrubbing determines whether the error is recoverable. If the error is recoverable Memory ProteXion technology will write the data to new location; if it is not recoverable, scrubbing sends an alert to light path diagnostics, which then notifies IBM Systems Director.

**Memory mirroring** works much like disk mirroring. The total memory is divided into two channels, a primary and a backup. Data is *written concurrently to both channels*. If one of the DIMMs in the primary channel fails, it is instantly disabled and the mirrored memory in the backup channel becomes active (primary) until the failing DIMM is replaced. With mirroring enabled, one-half of total memory is available for use.

Alternatively, with **memory rank sparing**, two *ranks* per memory card are configured as spares. These spare ranks are used in the event of a memory failure on the same memory card. This feature is more efficient than mirroring, reserving less memory and leaving more available for operating system use. (A memory card using dual-rank DIMMs, for example, would have up to 16 ranks, 14 available.)

Memory is available in **1GB, 2GB, 4GB, 8GB, 16GB or 32GB** DIMMs<sup>2</sup>.



## High-Performance Scalability

Memory scalability beyond the **1TB** capacity of the x3690 X5 chassis is achieved simply by cabling a 1U **MAX5** memory expansion unit to the chassis. It doubles memory capacity by adding another **32** DIMMs slots for up to **1TB** of additional memory (**2TB** total in only 3U of rack space).

The MAX5 scalability and memory expansion feature is a unique IBM enhancement, not offered by other x86 server architectures.

## High-Performance Adapter Slots

There are **five PCIe (PCI Express) Gen 2** adapter slots standard. **PCI Express** Gen 2 is the latest high-performance, low-latency, next-generation serial I/O bus. **Slot 1** is a **3/4-length/full-height** slot. **Slot 2** is **half-length/full-height**. **Slots 3 through 5** are **low-profile**. **Slots 1, 2, 3 and 5** are **x8** ("by 8") slots, both physically and electrically (meaning they operate at the full **x8 8GBps**<sup>9</sup> bidirectional speed). **Slot 4** is physically **x8** but electrically **x4**, which means it can accept a x8 adapter, but it will operate at x4 speeds (**4GBps**). If a ServeRAID controller is installed, it will go in **Slot 4**. If a **10Gb Ethernet Adapter** is installed, it will go in **Slot 5**. If you require a **x16** adapter slot, an optional riser card containing one **full-length/full-height x16/x16** slot can be added as a replacement for the riser card containing **Slots 1 and 2**.



---

Because the dual **Gigabit Ethernet**, **video**, and **systems management** controllers are all integrated onto the system board, four of the five adapter slots are *available*, which offers you a wide degree of latitude in expansion options.

---

## Storage Controllers

All models include an integrated **IBM ServeRAID-M1015 SAS/SATA** Controller. Several other controllers are supported, including:

- The **6Gbps (x8 PCIe) ServeRAID-M1015 SAS/SATA** controller — **RAID-0/1/10** (no cache) for up to **32** drives (limited to available bays). The **IBM ServeRAID M1000 Series Advance Feature Key** adds **RAID-5** and Self-Encrypting Disk (**SED**) support.
- The **6Gbps (x8 PCIe) ServeRAID-M5014 SAS/SATA** controller — enhanced performance with **256MB** of cache memory, and supports **RAID-0/1/10/5/50/6/60** for up to 32 drives (limited to available bays).
- The **6Gbps (x8 PCIe) ServeRAID-M5015 SAS/SATA** controller — enhanced performance with **512MB** of cache memory and battery backup, and supports **RAID-0/1/10/5/50/6/60** for up to 32 drives (limited to available bays).
- The **6Gbps (x8 PCIe) ServeRAID-M5025 SAS/SATA** controller — enhanced performance with **512MB** of cache memory and battery backup. It supports **RAID-0/1/10/5/50/6/60** for multiple IBM System Storage expansion units (with up to **240** HDDs).
- The **6Gbps (x8 PCIe) ServeRAID-B5015 SSD** controller — enhanced SSD performance with SSD-optimized microcode, offering the highest IOPS performance for SSDs in a RAID controller. It supports **RAID-0/1/5/6** for up to **8** drives.
- The **IBM 6Gbps SSD HBA** offers the highest IOPS performance for SSDs in a *non*-RAID controller.

The **ServeRAID M5000 Series Advance Feature Key** adds **RAID-6/60** with **SED** support to the M5014 and M5015. The **ServeRAID M5000 Series Battery Key** adds **battery backup** support to the M5014. The **ServeRAID M5000 Series Performance Accelerator Key** adds **SSD** support to the M5015.

Additional external SAS/SATA disk storage, as well as tape backup, is available via one of several supported iSCSI or FC SAN controllers.

---

## Drive Bays

x3690 X5 servers contain up to **16** HDD bays that support **2.5-inch hot-swap** internal **SAS/SATA** drives totaling up to **9.6TB**, or up to **24** solid-state drive bays, supporting **1.8-inch hot-swap SSDs** totaling up to **4.8TB**, using up to *three* IBM **eXFlash** storage options. Each individual SSD is capable of sustaining approximately **4,300 IOPS**, which is more than 10X what an enterprise 15,000 RPM 2.5-inch 146.8GB SAS HDD can achieve (420 IOPS). Thus, *each* eXFlash module (containing 8 SSDs) is capable of producing up to **34,400 IOPS**, vs. only **3,360 IOPS** for 8 SAS HDDs. This means one x3690 X5, with three eXFlash options (24 drives), is capable of generating **103,200 IOPS**, vs. only about 10,000 IOPS for typical storage servers using an equal number of HDDs. Combining 24 SSDs with 2 IBM 640GB High IOPS MLC Duo Adapters produces an astonishing **503,200 IOPS** for a single two-socket server.

Each customer-installable eXFlash option includes a drive cage, SSD backplane, and two x4 SAS/SATA cables, which can be installed in seconds as a single unit. For RAID support with optimum IOPS, use **one ServeRAID-B5015** controller per eXFlash option (or **two** for even more IOPS). For optimum *non-RAID* IOPS, use **one 6Gbps SSD HBA** controller per eXFlash option. (**Note:** Because SSDs offer triple<sup>4</sup> the reliability of SAS HDDs, RAID support is less necessary.) Each adapter requires a PCIe slot. For **maximum IOPS performance**, use **two ServeRAID-B5015** controllers per eXFlash option (maximum 5, due to available PCIe slots).

Supported single-chassis configurations include:

- **4 2.5-inch HDDs + 8 1.8-inch SSDs** (with 1 eXFlash option)
- **8 2.5-inch HDDs + 16 1.8-inch SSDs** (with 2 eXFlash options)
- **16 2.5-inch HDDs**
- **24 1.8-inch SSDs** (with 3 eXFlash options)
- **Plus 2 High IOPS adapters** (in 2 adapter slots)

For RAID support with HDDs, use one **ServeRAID-M1015**, **M5014**, or **M5015** controller per **8 internal drives**, or one **ServeRAID-M5025** controller or Fibre Channel adapter per multiple *external* IBM System Storage<sup>®</sup> expansion units (up to 240 drives).

An optional (ultraslim, 0.5") optical drive with a SATA interface is supported in all x3690 X5 servers. An external USB floppy drive or optical drive may be used, if needed.

For still more storage, NAS, iSCSI SAN, or FC SAN external expansion options can be added,

---



---

using an optional controller.

---

## Flexible Internal Storage Capacity

The eX5 servers offer a choice of internal storage, supporting up to **sixteen** high-capacity **2.5-inch** **hot-swap** Serial-Attach SCSI (**SAS**) or Serial ATA II (**SATA II**) hard disk drives, or up to **24** high-IOPS **1.8-inch** solid-state drives (**SSDs**) *chassis* (using three IBM eXFlash options). In addition, two **IBM High IOPS SSD adapters** can be used per node/chassis.

### 2.5-inch SAS

- **7,200 RPMs** — 160, **500GB (8TB maximum)**
- **10,000 RPMs** — 73.4, 146.8, 300, or **600GB (9.6TB)**
- **15,000 RPMs** — 73.4 or **146.8GB (2.38TB)**

2.5-inch drives not only require less space than 3.5-inch drives, they *weigh less, consume half the power, produce less noise, seek faster, and offer increased reliability.*

### 2.5-inch SATA

- **7,200 RPMs** — 160 or **500GB (8.0TB)**

### 1.8-inch SSD

- **50, 200GB (4.8TB maximum)**
  - **High I/O Performance**
    - Offers up to 8X more IOPS than HDDs (67/33% read/write OLTP transaction base mix);
    - Optimized for heavy mix of read and write operations, such as transaction processing, media streaming, surveillance, file copy, logging, backup / recovery, and business Intelligence
  - **Lower-Cost IOPS Performance**
    - Yields better \$/IOPS: lower capacity (GB) required to achieve higher IOPS
    - Consumes less energy and produces less heat than a conventional disk drive
  - **Superior Uptime**
    - 3X the reliability of mechanical disk drives
    - No moving parts to fail
    - Enterprise wear-leveling to extend life even further
  - **Full OS Support**
    - Supports all ServerProven® OSes

For read-intensive transactional workloads requiring extreme IOPS performance, such as database, video-on-demand, and caching, IBM offers **High IOPS SSD PCIe Adapters**, equivalent to the IOPS output of approximately **500<sup>5</sup>** 3.5-inch 600GB 15K SAS HDDs, with **99% lower latency** (30  $\mu$ s), and **7.7x** the bandwidth of a HDD. On a performance-per-watt basis, these adapters outperform HDDs by up to **445x<sup>6</sup>**.

### High IOPS SSD PCIe Adapters

- **IBM 160GB High IOPS SS Class SSD PCIe Adapter (x4 Gen 1—100K IOPS, 750MBps)**
- **IBM 320GB High IOPS SS Class SSD PCIe Adapter (x4 Gen 1—100K IOPS, 750MBps)**
- **IBM 320GB High IOPS MS Class SSD PCIe Adapter (x4 Gen 1—100K IOPS, 750MBps)**
- **IBM 320GB High IOPS SD Class SSD PCIe Adapter (x8 Gen 2—200K IOPS, 1.5GBps)**
- **IBM 640GB High IOPS MLC Duo Adapter (x8 Gen 2—200K IOPS, 1.5GBps)**

Because these adapters go in PCIe slots, rather than drive bays, they can be used *in addition* to standard HDDs and SSDs. This makes them ideal for virtualized servers running multiple tasks, requiring both write- and read-intensive storage media. As flash media they also offer a predictable lifetime, compared to the sporadic failure rate of HDDs, using N+1 chip-level redundancy and 11-bit ECC protection.

If you need more storage space, terabyte capacities are possible with external direct-attach NAS and SAN solutions.

---

## 10 Gigabit Ethernet Virtual Fabric Adapter

Selected x3690 X5 models include one **dual-port 10Gb Virtual Fabric** adapter. (It is optional for all other models.) The Emulex Virtual Fabric Adapter is an industry-leading performance and scalability per watt, dual-port network adapter for 10Gbps Ethernet (10GbE) networks. It offers the benefits and flexibility of I/O convergence in a single end-to-end solution. Protocol offload for stateless TCP/IP and TCP Chimney provide maximum bandwidth with minimum use of CPU





resources. It achieves line rate 10Gbps performance with support for TCP/IP stateless offloads and TCP Offload Engine (TOE) support. TOE reduces system processor utilization, providing increased system performance and reducing overall system power requirements.

The adapter is based on the Emulex OneConnect Universal Converged Network Adapter (UCNA) platform that also includes the capability for future upgrades to Fibre Channel over Ethernet (FCoE) and iSCSI protocol offloads. By using a common infrastructure for Ethernet and storage networks, data centers can reduce capital expense (CapEx) for adapters, switches and cables, and operational expense (OpEx) for power, cooling and IT administration.

End-to-end data protection with hardware parity, CRC, ECC and other advanced error checking and correcting ensure that data is safe from corruption.

#### Dual 10Gbps Ethernet ports:

- IPv4/IPv6 TCP, UDP checksum offload; Large Send Offload (LSO); Large Receive Offload; Receive Side Scaling (RSS); IPV4 TCP Chimney Offload
- VLAN insertion and extraction
- Jumbo frames up to 9000 Bytes
- Preboot eXecutive Environment (PXE) 2.0 network boot support
- Interrupt coalescing
- Load balancing and failover support including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support and IEEE 802.3ad.

**Note:** You must have either one SFP+ transceiver or one SFP+ direct-attached cable for *each* of the two 10Gb ports on the adapter.

## Dual Gigabit Ethernet Controller

The eX5 systems include one integrated **dual-port Broadcom 5709C** Gigabit Ethernet controller, with **TOE** (TCP Offload Engine) acceleration support, as well as **load-balancing** and **failover** capabilities between the two ports.

**TOE** support helps improve overall system performance by offloading TCP/IP protocol processing from the system microprocessor to the onboard Ethernet TOE processor.

The controller also supports highly secure remote power management using **IPMI 2.0**, plus **Wake on LAN®** and **PXE** (Preboot Execution Environment) flash interface. Optional PCI adapters offering failover and load balancing between adapters are available for added throughput and increased system availability.



## Integrated Virtualization

All models of the x3690 X5 support a **USB 2.0 Flash Key** installed preloaded with **VMware vSphere Hypervisor** (formerly ESXi). Rather than management through a Service Console based on a Linux operating system, vSphere Hypervisor relies on aggregate management tools, including VirtualCenter, the Remote Command Line interface and the introduction of CIM for standards-based and agentless hardware monitoring.

vSphere Hypervisor includes all the performance, scalability and compatibility features of a hypervisor installed on disk, including full **VMFS** support across FC SAN, iSCSI SAN, and NAS, and **4-way VSMP**. Because it runs from flash memory, it's extremely fast and ideal for diskless configurations. It also offers enhanced security, because it runs without an operating system-based console and is updated/patched much like firmware.




## Ultra-Efficient Cooling

**Five** strategically located **hot-swap/redundant** fans, combined with efficient airflow paths, provide highly effective system cooling for the eX5 systems, with N+1 redundancy.

The fans automatically adjust speeds in response to changing thermal requirements, depending on the zone, redundancy, and internal temperatures. When the temperature inside the server increases, the fans speed up to maintain the proper ambient temperature. When the temperature returns to a normal operating level, the fans return to their default speed. Why not simply run the fans at 100% capacity all the time? For several good reasons: to reduce the ambient noise, reduce the wear-and-tear on the fans and reduce the server power draw. The reduction in ambient noise and power draw may be relatively minor for a single server, but if you have a larger number of systems the aggregate savings in power and noise pressure can be significant.

In addition, the server uses **hexagonal ventilation holes** in the chassis. Hexagonal holes can be grouped more densely than round holes, providing greater airflow through the system cover.

This cooling scheme is important because newer, more powerful processors generate a significant amount of heat, and heat must be controlled for the system to function properly.



---

## Hot-Swap/Redundant Components

System availability is maximized through the extensive use of hot-swap and redundant components, including:

- **Hot-swap redundant hard disk drives** and **solid-state drives**
- **Hot-swap redundant cooling fans**
- **Hot-swap redundant power supplies**

---

## Light Path Diagnostics

Light path diagnostics enables a technician to quickly identify and locate a failed or failing system component, such as a specific memory DIMM. (Monitored components include processors, memory, power supplies, fans, HDDs, and adapters, plus system temperature.) This enables quick replacement of the component, which helps increase server uptime and lower operating costs.



The front of the server has an LED indicator light to show possible component failures. If the front LED indicates an error condition, by pressing a button on the front of the server an LED panel will pop out and drop down for easy viewing *without the need to open the server cover* or remove the server from the rack. The light path diagnostics panel tells the servicer which component requires attention. In addition, many components have their own identifying LEDs. For example, each of the memory modules has an LED next to the socket, as do all processors, all adapter slots, all fans, all power supplies, all voltage regulator modules and the service processor, allowing the servicer to easily identify exactly which component needs servicing. By following the “light path,” the component can be replaced quickly, and without guesswork. (**Note:** In the event of a failed DIMM, the system will restart and mark the DIMM as bad while offline, thus allowing the system to continue running, with reduced memory capacity, until serviced.)

In addition to the outstanding diagnostic detection, there is a new feature with a digital readout that indicates at what stage of the boot process your system may have experienced trouble. This can reduce the troubleshooting time and help get your system up and running quickly and efficiently.

---

## Energy-Smart Features

IBM servers work hard to keep your energy bills low—from high-efficiency power supplies and fans to lower-draw processors, memory, and solid-state drives. Technologies such as these, along with the Xeon E7 family and 7500/6500 series processors that intelligently adjust their voltage and frequency, help lower IT costs:

- **95W 8-core 7500 series** and **6-core E7 family processors** use **27%** less energy than 130W processors.
- **1.5V DDR3 DIMMs** consume **10-15%** less energy than 1.8V DDR2 DIMMs used in older servers.
- **1.35V DDR3 DIMMs** consume up to **19%** less energy than 1.5V DDR3 DIMMs.
- **Solid-State Drives** consume up to **80%** less energy than 2.5-inch HDDs and up to **88%** less than 3.5-inch HDDs.
- On a performance-per-watt basis, **High IOPS SSD Adapters** outperform enterprise HDDs by up to **445-to-1**<sup>7</sup>
- **Dynamic fan speeds** — In the event of a fan failure, the other fans will run faster to compensate until the failing fan is replaced. Competitive fans must run faster at all times, just in case, wasting power.

---

## Other Features

- **Trusted Platform Module standard** — This module provides a highly secure start-up process from power-on through hand-off to the operating system boot loader. ACPI support is provided to allow ACPI-enabled operating systems to access the security features of this module. (TCG V1.2-compliant.)
- **Eight USB 2.0 ports** — Provides flexibility to add high-speed external devices. The USB 2.0 specification supports up to 480Mbps transfer rates. (Note: Not all USB 2.0 devices are capable of achieving this rate.) **Two** ports are provided on the front of the server, **four** are on the back, and **two** are internal to support an optional flash memory key with embedded hypervisor.
- **Toolless slides** — Allows quick rack installation and quicker upgrade and servicing of the server.
- **Toolless chassis** — The cover can be opened without tools, and many components can be

added or replaced without tools, including processors<sup>8</sup>, DIMMs, hot-swap HDDs, PCIe adapters, an optional optical drive, and the integrated or optional ServeRAID controller. This can save a servicer significant time. **Note:** The chassis top cover should *not* be opened during server operation. Opening it will shut the server down.

---

## Advanced Systems Management Capabilities

The x3690 X5 has a high level of systems management capabilities that are well-suited to remote locations as well as to stand-alone environments. Features include UEFI, IMM with Virtual Media Key, IBM ToolsCenter, IBM Systems Director Active Energy Manager, Automatic Server Restart, Wake on LAN<sup>®</sup> support, PXE support, text console redirect, Predictive Failure Analysis, and IBM Systems Director.

The **IMM** provides industry-standard **Intelligent Platform Management Interface (IPMI) 2.0**-compliant systems management. It provides a number of important system functions, including:

- Monitoring of system and battery voltage, system temperature, fans, power supplies, processor and DIMM status
- Fan speed control
- Product ID and Family ID detection
- Highly secure remote power on/off
- System reset control
- NMI/SMI detection and generation
- System diagnostic LED control (power, HDD, activity, alerts, heartbeat)
- IPMI over LAN
- Serial Over LAN
- Proxy server support
- LAN messaging and alerting
- Text console redirection over LAN
- Predictive Failure Analysis for system fans
- Web-based out-of-band control
- SSL (Secure Socket Layer) and LDAP (Lightweight Directory Access Protocol) support
- VLAN support
- Enhanced authentication and encryption algorithms (RMCP+, AES)
- Local update of IMM firmware
- Firmware firewall
- Support for IPMI v2.0 compliant management software (e.g., xCAT)
- Other mandatory and optional IPMI IMM functions

The IMM alerts IBM Systems Director to anomalous environmental factors, such as voltage and thermal conditions—even if the server has failed.

The x3690 X5 also include standard an **IBM Virtual Media Key** for additional systems management capabilities, including:

- Latest OS failure screen capture
- Graphical console redirection over LAN
- Remote virtual floppy and CD-ROM
- High-speed remote redirection of PCI video, keyboard and mouse

**IBM ToolsCenter** consolidates 42 needed tools for managing servers individually into an integrated suite of 8 tools. They are organized by function: deployment, updates, configuration and diagnostics. Tools are now simpler to access and use with a single webpage for access, a common look and feel and a common command line interface for the scripting tools. The ToolsCenter **Bootable Media Creator** offers significantly more functionality than past tools with the ability to add more tools to the bootable image and to automatically download the bootable environment if needed. Bootable Media Creator allows you to create bootable CDs, DVD, and USB keys for updates customized to your systems.

**Automatic Server Restart (ASR)** helps reduce downtime by restarting the server automatically in the event of a system lockup. If the operating system crashes or the hardware freezes, it automatically triggers the ASR hardware, which immediately restarts the server (and logs an ASR event with IBM Systems Director). These features are designed so that *no more than five minutes can pass before the server is restarted*.

**Wake on LAN** permits the server to be remotely powered on if it has been shut off. Once

---

powered up, the server can be controlled across the network, using the **Preboot Execution Environment (PXE)**.

Like Wake on LAN, PXE is system firmware. It enables software such as the optional **IBM Remote Deployment Manager** to take control of a system before the BIOS, operating system or applications are loaded (using Wake on LAN/PXE) and lets an administrator perform many low-level tasks remotely that would otherwise require a visit to each system. These tasks may include such things as formatting a hard disk drive, updating system firmware, or deploying a Windows or Linux operating system.

**Text and Graphical Console Redirection** support allows the administrator to remotely view system text messages and graphics over serial or LAN.

**Predictive Failure Analysis (PFA)** is designed to allow the eX5 servers to detect impending failure of supported components (processors, memory, PCIe slots, power supplies, fans, the XCeL4v Dynamic Server Cache, and hard disk drives) *before* actual failure, and alert the administrator through IBM Systems Director. This gives you the ability to replace the failing component *before* it fails, resulting in increased uptime.

**IBM Systems Director** software for advanced workgroup management is included standard. IBM Systems Director comes with a portfolio of tools, including **IBM Systems Director Active Energy Manager**, **Service and Support Manager**, and others. *System Availability* (a no-charge download) and *Capacity Manager* (sold separately) are available as add-ons for additional server management and increased availability. IBM Systems Director provides a single uniform graphical interface for all of these systems management functions. IBM Systems Director enables you to customize thresholds and monitor system components (for things like temperature, voltage regulation, etc.) to help maximize uptime.

IBM developed **IBM Systems Director Active Energy Manager** to put control of system power-saving features at the fingertips of administrators. Active Energy Manager is designed to take advantage of new features, such as monitoring power usage and balancing the performance of the system according to available power input. It provides the ability to plan and predict power consumption based on your hardware configuration. It also helps enable you to reduce the infrastructure required for redundancy, by using fewer servers on smaller power feeds and potentially lowering your overall data center support costs. It does this by inventorying all components, then adding up the total power draw and tracking the usage. It also includes power management features to help administrators manage or reduce power usage.

---

## Extensive System Support Features

The IBM services and technical support portfolio provides world-class, consistent, high-quality service and support. The eX5 servers offer a number of tools and services designed to make ownership a positive experience. From the start, IBM programs make it easier for you to plan for, configure and purchase System servers, get them running and keep them running long-term. These features include IBM ServerProven, the IBM Standalone Solutions Configuration Tool, IBM System x and BladeCenter Power Configurator, IBM ServerGuide, IBM Systems Director Service and Support Manager, Product Customization Services and extensive technical support offerings.

The IBM **ServerProven** program provides the confidence that specific options and operating systems have been tested on the server and are officially supported to work together. It is updated frequently to ensure that the latest compatibility information is always at your fingertips.

The IBM **Standalone Solutions Configuration Tool (SSCT)** is a downloadable tool that simplifies the often complex chore of configuring a full rack of servers (including blade servers) and confirming that you have all the cables, power distribution units, KVM (keyboard, video and mouse) switch boxes and other components you need, as well as the proper airflow clearances, electrical circuits and other environmental conditions.

**IBM System x and BladeCenter Power Configurator** helps IT managers plan for data center power needs, by providing the following information for specific configurations of System x and BladeCenter systems: *power input* (watts), *PDU sizing* (amps), *heat output* (BTUs), *airflow requirements through chassis* (CFM), *VA rating*, *leakage current* (mA), and *peak inrush current* (amps).

**IBM ServerGuide** (installed from CD) simplifies the process of installing and configuring System x and BladeCenter servers. ServerGuide goes beyond mere hardware configuration by assisting with the automated installation of the Microsoft® Windows® Server 2008 operating systems, device drivers and other system components, with minimal user intervention. (Drivers are also included for support of Novell NetWare, Red Hat Linux and SUSE LINUX.) This focus on deployment helps you reduce both your total cost of ownership and the complexity that administrators and technical personnel face.

**IBM Systems Director Service and Support Manager** (previously called IBM Electronic Service Agent™) is an innovative “call home” feature that allows System x and BladeCenter servers to automatically report hardware problems to IBM support, which can even dispatch

---



onsite service if necessary to those customers entitled to onsite support under the terms of their warranty or an IBM Maintenance Agreement. Service and Support Manager resides on a server and provides electronic support and problem management capabilities through a highly secure electronic dialogue between your systems and IBM. It monitors networked servers for hardware errors and it can perform hardware and software inventories and report inventory changes to IBM. All information sent to IBM is stored in a highly secure database and used for improved problem determination.

Additional services include hardware warranty upgrades and factory-installed **Product Customization Services** (PCS), such as asset tagging, hardware integration, software imaging and operating systems personalization.

IBM offers extensive industry-leading **technical support** by phone and via the Web. Support options include links to forums/newsgroups, problem submission, online shopping support, service offerings, device drivers for all IBM product lines, software downloads and even upcoming technical seminar worldwide schedules and registration. Also available are remote installation, configuration and usage support for System x hardware and software, as well as onsite custom services to provide the level of expertise you require.

**IBM Maintenance and Technical Support** solutions can help you get the most out of your IT investment by reducing support costs, increasing availability and simplifying management with integrated support for your multiproduct, multivendor hardware and software environment. For more information on hardware maintenance, software support, solution support and managed support, visit <http://ibm.com/services/maintenance>.

---

## Key Options

### IBM options for System x servers help you take your servers to a higher level

You can rely on System x options to supply a complete solution for your business needs. Options help you create an optimized server system to meet your data protection, storage and availability needs. Every IBM option is designed and tested for peak performance and flexibility, helping to maximize your return on investment. The combination of System x servers and options lets you keep your fingers on the pulse of your e-business.

**Processors** — The Intel Xeon processors provide high clock rates, a choice of 10, 8, 6, or 4 cores, a large L3 cache and advanced features for availability and manageability (processor-specific). Large cache size, combined with fast **1066MHz** memory access, reduces memory latency and facilitates the movement of data through the processor and I/O devices. (**Note:** System performance depends not only on the number of processors in the server but also on the power and functionality of each processor.) Adding a second processor may be a cost-effective way to achieve significant performance improvements.

**Memory** — Memory is a significant factor in systems application performance. Adding more memory to a System x server is one of the most effective ways to increase application performance. For optimum performance, an 8-core processor should have twice as much memory as a 4-core processor. x3690 X5 servers support **32 DIMMs** natively and **64 DIMMs** using a **MAX5** memory expansion unit. **1.35V DIMMs** (machine type 7147) can help reduce energy consumption by up to **19%**.

**Hard Disk Drives** — IBM hard disk drives help customers improve the transaction and cost performance of their System x servers. The choice of hard disk drives can be a critical aspect of maximizing the I/O throughput of the system. **SAS** hard disk drives (2.5-inch) are available for the eX5 servers with capacities of **500GB** at **7,200 RPMs**, up to **600GB** at **10,000 RPMs**, and **146.8GB** at **15,000 RPMs**. Enterprise-class 2.5-inch **SATA** hard disk drives are available with capacities of up to **500GB** at **7,200 RPMs**.

**Solid State Drives** — IBM offers a choice of **50GB** or **200GB** 1.8" **solid-state drives** as a higher reliability, lower-energy-use, higher-IOPS alternative to internal HDDs. They can be used as a highly available boot drive, for storing disk images, or for other uses that stress read performance, such as databases.

**High IOPS SSD PCIe Adapters** — For storage media offering up to **500 times** the I/O operations per second (IOPS) and up to **7.7 times** the bandwidth of 15,000-RPM HDDs, IBM offers the **160GB High IOPS SS Class SSD PCIe Adapter**, the **320GB High IOPS SS Class SSD PCIe Adapter**, the **320GB High IOPS MS Class SSD PCIe Adapter**, and the **320GB High IOPS SD Class SSD PCIe Adapter**, as well as the **640GB High IOPS MLC Duo Adapter**. These adapters offer near-DRAM performance, with extremely high data retention (up to 25 years) and RAID-grade data protection, and can be used in *addition* to HDDs and SSDs.

**IBM ServeRAID Controllers** — System x servers using ServeRAID technology allow companies to build a reliable foundation for business-critical computing. IBM ServeRAID technology allows an array consisting of multiple physical hard disk drives to be treated as one logical drive. ServeRAID technology also allows data to be stored redundantly, across multiple hard disk drives—enhancing both the integrity and the availability of the data. SAS and SATA ServeRAID controllers offer enhanced performance due to onboard processors and cache. Because IBM




ServeRAID controllers can help significantly improve data transfer rates, this technology is extremely effective when implementing demanding, transaction-oriented applications. By employing the advanced fault tolerance of IBM ServeRAID technology, companies can effectively implement networked business systems that require large amounts of storage space for data and applications that must be available for their businesses to continue operating.

Supported SAS/SATA RAID controllers include the IBM **ServeRAID-M1015** (x8 PCIe, RAID-0/1/10), IBM **ServeRAID-M5014** (x8 PCIe, **256MB** cache, RAID-0/1/10/5/6; optional battery backup, RAID-50/60 and SED), IBM **ServeRAID-M5015** (x8 PCIe, **512MB** cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED) and the IBM **ServeRAID-M5025** (x8 PCIe, **512MB** cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED; supports multiple external IBM **System Storage** expansion units). The **ServeRAID-B5015** controller is optimized for high-IOPS RAID performance with SSDs. The IBM **6Gbps SSD HBA** is optimized for high-IOPS *non*-RAID performance.


**External Storage** — The IBM **System Storage EXP810**, **EXP3512** and **EXP3524** expansion units, as well as the **DS3000**, **DS4000**, **DS5000**, **DS8000** series, and **XIV** storage subsystems and **N3000**, **N5000**, **N6000**, and **N7000** NAS systems, comprise a powerful and broad shared storage family with integrated management software designed to meet midrange and enterprise needs. In addition, IBM **Storwize® V7000** is virtual storage that offers high efficiency and flexibility through built-in SSD optimization and “thin provisioning” technologies.

## **x3690 X5 Images**

### **Front View**




### **Rear View**




*Ultrascaleability in a 2-socket package, combined with extreme availability features and industry-leading performance*

### **Interior View**



### **MAX5 Interior View**



Ultrascaleability in a 2-socket package, combined with extreme availability features and industry-leading performance

x3690 X5 Specifications				
Machine type	7147-A1x/A2x/A3x/A5x/A6x/A7x, C1x, D1x/D2x. H2x/H3x 7148-1Rx, 2Dx/2Rx, 3Dx/3Gx/3Rx/3Sx, 4Rx/4Sx, ARx			
Form factor (standard/maximum)	2U / 3U (1 2U chassis + 1 1U MAX5 memory expansion unit)			
Processor type	Ten-core Xeon (E7-28xx, E7-88xx) 2.0GHz E7-2850 (A5x), 2.13GHz E7-8867L (CTO), 2.26GHz E7-2860 (A6x, D1x/D2x), 2.4GHz E7-2870 (A7x, H3x/H3x), 2.67GHz E7-8837 (C1x)	Eight-core Xeon (E7-28xx/X65xx/X75xx) 2.0GHz X6550 (3Dx/3Gx/3Rx), 2.0GHz X7550 (3Sx), 2.0GHz E7-2820 (A2x), 2.13GHz E7-2830 (A3x), 2.26GHz X7560 (4Rx/4Sx)	Six-core Xeon (E7-28xx/E7-48xx/E65xx) 1.73GHz E7-2807 (A1x), 1.86GHz E7-4807 (CTO), 2.0GHz E6540 (2Dx/2Rx)	Four-core Xeon (E75xx) 1.86GHz E7520 (1Rx, ARx)
Processor cores	20 (A5x/A6x/A7x, C1x, plus E7-8867L, via CTO)	16 (3Dx/3Gx/3Rx/3Sx, 4Rx/4Sx, A2x/A3x)	12 (2Rx, A1x, plus E7-4807, via CTO)	8 (1Rx, ARx)
Maximum processor power draw	130W (3Dx/3Gx/3Rx/3Sx, 4Rx/4Sx, A5x/A6x/A7x, C1x, D1x/D2x, H2x/H3x)	105W (1Rx, 2Dx/2Rx, A1x/A2x/A3x/ARx, plus E7-8867L, via CTO)		95W (E7-4807, via CTO)
QuickPath Interconnect (QPI) speed (gigatransfers per second)	6.4 GTps (2Dx/2Rx, 3Dx/3Gx/3Rx/3Sx, 4Rx/4Sx, A3x/A5x/A6x/A7x, C1x, D1x/D2x, H2x/H3x, plus E7-8867L, via CTO)	5.86 GTps (A2x)		4.8 GTps (1Rx, A1x/ARx plus E7-4807, via CTO)
# of processors standard / maximum	2 / 2 (2Dx, 3Dx, 4Sx, D1x/D2x, H2x/H3x)		1 / 2 (1Rx, 3Sx, 4Rx, A1x/A2x/A3x/A5x/A6x/A7x/Arx, C1x)	
Internal L3 cache	30MB (A7x, H2x/H3x, plus E7-8867L, via CTO)	24MB (4Rx, 4Sx, A3x/A5x/A6x, C1x, D1x/D2x)	18MB (2Dx/2Rx, 3Dx/3Gx/3Rx/3Sx, A1x/A2x plus E7-4807, via CTO)	12MB (1Rx, ARx)
Chipset	Intel 7500 & IBM eX5			
Maximum memory / with MAX5	1TB / 2TB			
Standard memory <sup>9</sup>	256GB (32 x 4GB + 32 x 4GB in MAX5)—2Dx	192GB (16 x 16GB)—H2x/H3x	64GB (16 x 4GB)—D1x/D2x	16GB (4 x 4GB)—3Dx
	12GB (4 x 2GB + 2 x 2GB in MAX5)—3Sx, 4Sx	8GB (4 x 2GB)—1Rx, 2Rx, 3Rx/3Gx, 4Rx	8GB (2 x 4GB)—A1x/A2x/A3x/A5x/A6x/A7x, C1x	4GB (2 x 2GB)—ARx
# of DIMM sockets total / available	64 / 0 (32 / 0: chassis, 32 / 0: MAX5)—2Dx		64 / 58 (32 / 28: chassis; 32 / 30: MAX5)—3Sx, 4Sx	
	32 / 28 (1Rx, 2Rx, 3Dx/3Rx/3Gx, 4Rx)		32 / 30 (A2x/A3x/A5x/A6x/A7x/ARx, C1x)	
MAX5 memory expansion unit capacity	1TB (32 x 32GB)			
Maximum memory with 1 server and 1 MAX5 expansion unit	2TB (64 x 32GB)			
Supported DIMMs	PC3-10600 — 1333MHz 2GB, 1Rx8, 1.5V 2GB, 2Rx8, 1.5V 4GB, 2Rx8, 1.5V	PC3L-10600R-999 — 1333MHz 4GB, 2Rx8, 1.35V	PC3-8500 — 1066MHz 4GB, 4Rx8, 1.5V 8GB, 4Rx8, 1.5V 16GB, 4Rx4, 1.5V	PC3L-8500 — 1066MHz 8GB, 4Rx8, 1.35V 16GB, 4Rx4, 1.35V 32GB, 4Rx4,



Ultrascaleability in a 2-socket package, combined with extreme availability features and industry-leading performance

x3690 X5 Specifications				
				1.35V <sup>10</sup>
Memory interleaving	Yes (two-way using pairs of DIMMs)			
Chipkill protection supported	Yes (using x4 or x8 DIMMs)			
Memory ProteXion supported	Yes			
Memory rank sparing supported	Yes			
Memory mirroring supported	Yes			
Memory scrubbing supported	Yes			
# of HDD/SSD drive bays maximum / available	4 / 4 2.5-inch HDDs + 8 / 8 1.8-inch SSDs (using 1 optional eXFlash pack); or 8 / 8 2.5-inch HDDs + 16 / 16 1.8-inch SSDs (using 2 optional eXFlash packs); or 16 / 16 2.5-inch HDDs; or 24 / 24 1.8-inch SSDs (using 3 optional eXFlash packs)			
# of HDD/SSD drive bays standard	1 x 2.5-inch HDD backplane (4 bays) and 1 x 1.8-inch SSD backplane (8 drives) included—3Dx (2 additional 1.8-inch SSD backplanes available via eXFlash packs)		1 x 2.5-inch HDD backplane (4 bays) included (all other models)—3 additional 1.8-inch SSD backplanes available via eXFlash packs	
	2 x 2.5-inch HDD backplanes (8 bays) + 2 x 2.5-inch HDD backplanes (8 bays) included—H2x (1 additional 1.8-inch SSD backplane available via eXFlash packs)		2 x 1.8-inch SSD backplanes (16 bays) included—D1x/H3x (1 additional 1.8-inch SSD backplane available via eXFlash pack)	
# of 5.25" bays total / available	1 / 1 (optional optical drive with SATA interface)			
Maximum HDD/SSD capacity	2.5-inch SAS 9.6TB (16 x 600GB) hot-swap		2.5-inch SATA 8TB (16 x 500GB) hot-swap	1.8-inch SSD 4.8TB (24 x 200GB) hot-swap
Internal HDD/SSD capacities supported	2.5-inch SAS 500GB — 7.2K RPMs; 146.8, 300, 600GB — 10K RPMs; 73.4, 146.8GB — 15K RPMs		2.5-inch SATA 160, 500GB — 7.2K RPMs;	1.8-inch SSD 50,200GB
# of HDDs/SSDs included	16 x 200GB SSDs (D1x)	10 x 200GB SSDs (H3x)	8 x 50GB SSDs + 8 x 300GB HDDs (H2x)	None (all other models)
Additional Internal SSD storage via PCIe adapter	160GB High IOPS SS Class SSD PCIe Adapter 320GB High IOPS SS Class SSD PCIe Adapter 320GB High IOPS SD Class SSD PCIe Adapter 320GB High IOPS MS Class SSD PCIe Adapter 640GB High IOPS MLC Duo Adapter			
# of optical drives standard	None (optional optical drive in dedicated 5.25" UltraBay)			
Storage technology	Hot-swap SAS/SATA/SSD			
External disk drives supported	Yes, via optional ServeRAID-M5025 controller and Fibre Channel adapters			
External disk storage devices supported	System Storage EXP810, EXP3512, EXP3524 expansion units, DS4200, DS4700, DS4800, DS5100, DS5300, DS8100, DS8300, N3300, N3600, N3700, N5200, N5300, N5500, N5600, N6040, N6060, N6070, N7600, N7700, N7800, N7900 storage subsystems, plus XIV and Storwize V7000			
Integrated RAID controller / cache	ServeRAID-B5015 (x8 PCIe, no cache, RAID-0/1/5 for SSD)—D1x	ServeRAID-M1015 (x8 PCIe, RAID-0/1/10 for HDD/SSD)—all other models	IBM 6Gbps SAS HBA (x8 PCIe, non-RAID controller for SSD)—D2x	

Ultrascaleability in a 2-socket package, combined with extreme availability features and industry-leading performance

x3690 X5 Specifications			
Optional RAID controllers supported	ServeRAID-B5015—D2x	ServeRAID-M5015 (x8 PCIe, 512MB cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED)—all HDD models	ServeRAID-M5025 (x8 PCIe, 512MB cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED; support for multiple external IBM System Storage expansion units—all models)
	ServeRAID-M5014 (x8 PCIe, 256MB cache, RAID-0/1/10/5/6; optional battery backup, RAID-50/60 and SED)—all HDD models		
Embedded USB Flash Key	Yes (2Dx)		Optional (all other models)
# of adapter slots total / available	5 / 5 (ARx)	5 / 3 (2DX, 3Dx/3Gx, D1x/D2x, H2x/H3x)	5 / 4 (all other models)
# of PCIe physical x16/electrical x16 Gen 2 slots (16GBps)	1 optional full-height/full-length (requires replacing riser card containing two x8 slots)		
# of PCIe physical x8/electrical x8 Gen 2 slots (8GBps)	1 3/4-length/full-height / 1 half-length/full-height / 2 low-profile		
# of PCIe physical x8/electrical x4 Gen 2 slots (4GBps)	1 low-profile (used by RAID controller)		
# of video ports	2 (front and rear)		
Video controller	Matrox G200eV graphic controller integrated in IMM (optional nVIDIA Quadro FX 3800 graphics card supported)		
Video memory	16MB DDR2 SDRAM standard (shared with IMM)		
Maximum video resolution at 32-bit color	1024x768 resolution (analog), with a color depth of 32 bits at 85Hz; 1440x90 resolution (digital), with a color depth of 32 bits at 60Hz		
10 Gigabit Ethernet Virtual Fabric Adapter	Dual-port Emulex controller standard—in adapter Slot 5 (2DX, 3Dx/3Gx, D1x/D2x, H2x/H3x)	Optional (all other models)	
Gigabit Ethernet controller	Dual-port Broadcom BCM5709C		
10 Gigabit Ethernet controller	Emulex 10Gbps Virtual Fabric Adapter (standard or optional—model-specific)		
Network acceleration	TOE / iSCSI / RDMA (1Gb and 10Gb)		
Failover / load-balancing-capable	Yes / Yes		
# of Gigabit Ethernet ports	2 (rear) + 1 100Mb (rear) for systems management		
# of serial ports	1 (rear)		
# of USB 2.0 ports	6 external (2 front, 4 rear) ports, plus 2 internal USB connectors for hypervisor		
Integrated systems management controller	IMM		
Light path diagnostics support	Processors, memory, power supplies, fans, HDDs, adapter slots, and system temperature, with external pop-out/drop-down panel		
Predictive Failure Analysis support	Processors, memory, PCIe slots, power supplies, fans, hard disk drives, and solid-state drives		
Power supply size (chassis / MAX5)	675W universal, autoswitching, hot-swap (rear access)—chassis	675W universal, autoswitching, hot-swap (rear access)—MAX5	
# of power supplies standard / maximum (chassis / MAX5)	2 / 4 (chassis)	1 / 2 (MAX5)	
Hot-swap/redundant power supported (chassis / MAX5)	Yes / Yes		
# of fans/blowers standard /	5 / 5 (chassis)	5 / 5 (MAX5)	



*Ultrascaleability in a 2-socket package, combined with extreme availability features and industry-leading performance*

<b>x3690 X5 Specifications</b>		
<b>maximum</b>		
<b>Hot-swap/redundant fans supported (chassis)</b>	Yes	
<b>Rack mount method</b>	Slides and Cable Management Arm (provided standard)	
<b>Maximum altitude</b>	7,000 ft; 2,133 m	
<b>Heat emitted: maximum (chassis / MAX5)</b>	5115 <b>BTUs</b> ; 1498 <b>Watts</b> (x3690 X5)	??? <b>BTUs</b> ; ??? <b>Watts</b> (MAX5)
<b>Dimensions (chassis / MAX5)</b>	6.8" (172.8mm) <b>H</b> 17.5" (444mm) <b>W</b> 28.04" (712mm) <b>D</b>	1.75" (44.4mm) <b>H</b> 17.5" (444mm) <b>W</b> 28.35" (720.2mm) <b>D</b>
<b>Weight (chassis / MAX5)</b>	77 – 110 <b>lbs</b> (maximum) 35.3 – 49.9 <b>kg</b>	28.2 – 33.8 <b>lb</b> (maximum) 12.8 – 15.1 <b>kg</b>
<b>Operating temperature range</b>	50–95° F; 10–35° C (up to 3,000 ft / 914.4 m)	50–90° F; 10–32° C (3,000 ft to 7,000 ft / 914.4m to 2,133m)
<b>Operating systems supported</b>	Microsoft Windows Server 2008 / R2 (Standard/Web/Enterprise/Datacenter Editions) 64-bit; Windows HPC Server 2008; Windows 2008 Hyper-V; Microsoft Windows Server 2003 / R2 (Standard/Web/Enterprise/Datacenter Editions) 32/64-bit; RHEL 5/6 Server 64-bit (with and without Xen); SLES 10/11 64-bit (with and without Xen); VMware vSphere 4.0 / 4.1	
<b>Length of limited warranty</b>	3 years (parts and labor) <sup>11</sup>	

## ***The Bottom Line***

The eX5 servers are extremely powerful systems, incorporating leading-edge industry-standard features and adding IBM-unique innovations:

### **Price/Performance**

- **High-throughput processors** — 2.0 to 2.67GHz 10-core, 2.0 to 2.26GHz 8-core, 1.73 to 2.0GHz 6-core, or 1.86GHz 4-core Xeon processors; up to 20 (10-core), 16 (8-core), 12 (6-core), or 8 (4-core) threads per processor
- **Large L3 cache** — 30MB, 24MB, 18MB, or 12MB of L3 processor cache (model-specific)
- **Fast memory access** — Up to 1066MHz (processor-specific); **two-way interleaving** and a **low-latency memory controller**
- **Energy-efficient memory** — 1.35V DIMMs supported in 7147 machine type. (Uses low-power memory buffers as well.)
- **Embedded virtualization** standard or optional
- **Fast disk technology** — SAS/SATA RAID controller and drives standard (model dependent)
- **Fast SSD technology** — Optional SSD-optimized controller and drives
- **High-IOPS technology** — Optional **High IOPS** SSD adapters
- **Fast communications** — **Optional 10GbE Virtual Fabric Adapter** and integrated **dual-port Gigabit Ethernet** controller, supporting **load-balancing** and **failover**, as well as **TOE** acceleration
- **Fast I/O** — Four **x8 (8Gbps) PCIe Gen 2** and one **x4 (4Gbps) Gen 2** adapter slots

### **Flexibility**

- **Scalable** — From 4 to 20 processor cores, and 32 to 64 DIMMs, scalable on demand (using 1 MAX5 and 1 chassis)
- **Huge memory capacity** — Up to 1TB of DDR3 memory (model-specific), with 32 DIMM slots, plus another 1TB in 32 DIMM slots in a MAX5 expansion unit
- Integrated **flash** drive, preloaded with a hypervisor for "instant" virtualization (standard on selected x3690 X5 models, optional on others)
- A choice of up to 16 2.5-inch hot-swap SAS/SATA HDDs or 24 1.8-inch hot-swap SSDs internally, or a combination of the two)
- Optional IBM 160GB, 320GB, or 640GB High IOPS SSD Adapters in addition to HDDs and SSDs, for extremely high IOPS and bandwidth storage




































































































































- **High-performance external I/O expansion** — **Eight** 480Mbps **USB 2.0** ports (two front, four rear, two internal)
- Hardware-based **RAID-0/1/10** support standard in all models; optional RAID support for **RAID-10/5/50/6/60**
- **Five PCIe** slots (4 x 8GBps, 1 x 4GBps)
- Optical drive

**Manageability, Serviceability and Availability**

- **IBM Systems Director** systems management software, including:
  - IBM Systems Director Active Energy Manager
  - IBM Service and Support Manager
- **Integrated Management Module (IMM):**
  - ☐ **IPMI 2.0** compliance, including highly secure remote power control
  - ☐ **Text console redirection** systems management standard
  - ☐ Supports **LDAP** and **SSL** industry standards
- **Active Memory protection:**
  - ☐ Advanced **Chipkill** ECC memory correction
  - ☐ **Memory ProteXion**
  - ☐ **Spare rank memory**
  - ☐ **Memory mirroring**
  - ☐ **Memory scrubbing**
- Hardware-based **RAID-1/10** disk mirroring standard in a dedicated slot; optional **RAID-10/5/50/6/60** highly available arrays
- **Ultra-efficient cooling**
- **Hot-swap HDDs and/or SSDs**
- **Hot-swap/redundant power supplies**
- **Hot-swap/redundant cooling**
- **Light path diagnostics** (front LED panel, drop-down light path panel) monitoring processors, memory, power supplies, fans, HDDs, PCIe adapter slots, and system temperature
- **PFA support** for processors, memory, PCIe slots, power supplies, fans, hard disk drives, and solid-state drives
- **Toolless chassis** and **toolless slide** design; integrated **Cable Management Arm** simplify installation and removal of the server from the rack

## Server Comparison Chart

The following table shows the suggested uses for the respective IBM System x rack-optimized servers, including comparisons of the uses for which each server is best suited:

		Requirements										Rack-Optimized Servers											
Theme	Key Workloads	Scalability	Floating Point Performance	Memory Throughput	Integer Performance	I/O and Storage	Density	High Availability	Systems Management	Security	Distributed Deployment												
HPC	Cluster / HPC																						
	Modeling & Simulation																						
	High Performance DB																						
	Business Intelligence																						
Web 2.0 / Web 3D	Search																						
	Content																						
	Communities																						
	Commerce																						
Business Applications	Collaboration																						
	ERP/SCM																						
	CRM																						
	Hosted Client																						
Infrastructure Applications	Point of Sale																						
	Branch Office																						
	Virtualization																						
	Business Continuity																						
	Database																						
	Email/Collaboration																						
	Security																						
	Web Serving																						
	File & Print																						

Important  
  
 Nice to Have  
  
 Can do without  
  
 Best  
  
 Better  
  
 Good  


## For More Information

IBM System x Servers	<a href="http://ibm.com/systems/x">http://ibm.com/systems/x</a>
IBM Systems Director Service and Support Manager	<a href="http://ibm.com/support/electronic">http://ibm.com/support/electronic</a>
IBM System x and BladeCenter Power Configurator	<a href="http://ibm.com/systems/bladeCenter/resources/powerconfig.html">http://ibm.com/systems/bladeCenter/resources/powerconfig.html</a>
IBM Standalone Solutions Configuration Tool	<a href="http://ibm.com/systems/x/hardware/configtools.html">http://ibm.com/systems/x/hardware/configtools.html</a>
IBM Configuration and Options Guide	<a href="http://ibm.com/systems/x/hardware/configtools.html">http://ibm.com/systems/x/hardware/configtools.html</a>
IBM ServerProven Program	<a href="http://ibm.com/systems/info/x86servers/serverproven/compat/us">http://ibm.com/systems/info/x86servers/serverproven/compat/us</a>
Technical Support	<a href="http://ibm.com/server/support">http://ibm.com/server/support</a>
Other Technical Support Resources	<a href="http://ibm.com/systems/support">http://ibm.com/systems/support</a>

## Legal Information

© IBM Corporation 2010

IBM Systems and Technology Group  
Dept. U2SA  
3039 Cornwallis Road  
Research Triangle Park, NC 27709

Produced in the USA  
April 2011  
All rights reserved

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven®. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, ibm.com, Active Energy Manager, Active Memory, Chipkill, DB2, Memory ProteXion, OnForever, ServeRAID, ServerGuide, System Storage, Storwize, System x, Wake on LAN, and XpandOnDemand are trademarks of IBM Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available at <http://ibm.com/legal/copytrade.shtml>.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, the Windows logo, Hyper-V, SQL Server, and Windows are trademarks or registered trademarks of Microsoft Corporation.

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

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM

intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND 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 publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.

XSO03120-USEN-02