

Optimized for virtualization and database applications with maximum memory and compute capacity in a blade



Product Guide

October 2010

IBM BladeCenter HX5

Product Overview

Scalable 2-to-4-socket blade server optimized for virtualization/consolidation, database, and ERP

Suggested uses: Front-end and mid-tier applications requiring high performance (CPU, memory or I/O), enterprise-class availability, and extreme flexibility and power efficiency.

CONTENTS

Product Overview	1
Selling Features	2
Key Features	5
Key Options	13
HX5 Images	15
HX5 Specifications	16
The Bottom Line	18
Server Comparison	20
For More Information	21
Legal Information	21

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

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 servers to run business critical applications, the systems that run these applications need the type of technology designed into IBM's eX5 technology systems. The eX5 chipset represents a \$100M+ investment in designing a flagship offering that can harness the power of 4-socket-and-up 64-bit x86 (x64) Xeon processors. The eX5 family includes a scalable performance **blade server** with the ability to scale from a **single-wide blade** (30mm) to a **double-wide blade** (60mm). Maybe you'd like to start out with a 2-socket blade and possibly add sockets later, if your needs change. Or perhaps you need more than a 2-socket blade, but don't want to get locked into a monolithic 4-socket blade—again, in case your requirements change. With IBM, you can start at 2 sockets and grow to 4 if needed. And if you require a variety of 2- and 4-socket servers in your data center, you only have to qualify one server for all these workloads. This can save much time and effort and speed deployment. You can also save money on software licensing by virtualizing a 4-processor server into many VMs, rather than using multiple 2-processor servers. Huge amounts of memory also enable more or larger VMs, and larger databases (especially databases stored entirely in memory).

Reducing an entire server into a little over .5U of rack space (i.e., up to 14 servers in 9U) does *not* mean trading away features and capabilities for smaller size. The **IBM BladeCenter® HX5** blade server offers features comparable to many 1U rack-optimized full-featured servers, and then some: The HX5 supports up to **two** of the latest high-performance **8-core, 6-core, or 4-core** Intel Xeon **7500** and **6500** Series processors. The Xeon processors are designed with up to **24MB** of shared cache and leading-edge memory performance (up to **1066MHz** using **MAX5**, depending on processor model) to help provide the computing power you require to match your business needs and growth. In addition, the HX5 supports up to **128B** of **registered** double data rate III (**DDR3**) **ECC** (Error Checking and Correcting) memory in **16** DIMM slots, with multiple levels of IBM Active Memory™ protection, for the highest levels of performance and availability. Active Memory protection includes **IBM Memory ProteXion™**, **IBM Chipkill™** memory, **memory scrubbing**, **memory rank sparing**, and **memory mirroring**.

Because business requirements change and a 2-socket server that meets those needs today may not meet them in the future, the HX5 was designed to be upgradeable to meet the diverse needs of multiple workloads. For **compute-intensive workloads**, it can be configured as a 2-wide blade server with up to **4 processors / 32 cores, 32 DIMMs (256GB), 4 PCIe cards, 16 I/O ports**, and **4 solid-state drives (SSDs)**. Conversely, for **memory-intensive workloads**, the HX5 can be configured as a 2-wide server consisting of **1 blade server and 1 MAX5** memory expansion blades, with up to **40 DIMMs (320GB)**. **IBM FlexNode partitioning** allows a physical 4-socket configuration to be remotely reconfigured by software into *two* logical 2-socket servers, as needs change or for daily peak and off-peak workloads.

The HX5 supports **VMware vSphere Hypervisor** (formerly ESXi) preloaded on a standard USB flash drive. 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'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.

An integrated **dual-port Gigabit Ethernet¹** controller is standard, providing high-speed data transfers and offering **TOE** (TCP Offload Engine) support, **load-balancing** and **failover** capabilities. Via optional expansion cards, each blade can also connect to additional Ethernet, **10Gb Ethernet**, **SAS**, **Fibre Channel**, **iSCSI**, **InfiniBand™** and other high-speed communication switches housed in the chassis. This blade is designed with **power management capability** to provide the maximum uptime possible for your systems. In extended thermal conditions or power brownouts, rather than shut down completely, or fail, the HX5 automatically reduces the processor frequency to maintain acceptable thermal and power levels.

All HX5 models offer impressive features at an equally impressive price, including up to **two 1.8-inch fixed solid-state drives (SSDs)** with **RAID-0/1** support, and one optional **internal USB flash drive** (for embedded hypervisor). Additional direct-attach storage is available via the BladeCenter S chassis. Moreover, the HX5 is **optimized** for diskless operation, offering each blade server access to essentially unlimited external storage capacity via Fibre Channel, SAS, or iSCSI.

A single **BladeCenter H** chassis supports up to **14 hot-swappable** 30mm-wide HX5 blades in only **9U** (BladeCenter H) of rack space, or up to **12** in the **12U BladeCenter HT** high-speed telecommunications chassis. In addition to the blade servers, these chassis also hold up to **10** (BladeCenter H) or **8 switches/bridges** (BladeCenter HT) internally. The **BladeCenter S**, designed for SMB and mid-market customers, takes integration and affordability to a new level, combining up to **12** hot-swap **SAS/SATA** HDDs (with optional SAS card) and up to **6 blade servers** and **4 switches**. Not only can this save significant data center space (and therefore the cost of floor space and rack hardware) compared to 1U servers, it also consolidates switches/bridges and cables for reduced complexity and lower cabling costs, and it allows clients to manage everything in the solution as one. Using a BladeCenter H chassis, up to **56 HX5** servers (**112 processors/896 cores**) can be installed in one **industry-standard 42U** rack; but the value of BladeCenter extends far beyond high density data center environments.

The various BladeCenter chassis are designed to monitor environmental conditions in the chassis and each blade and send alerts to the administrator. Advanced standard features, such as **Active Memory**, **Predictive Failure Analysis**, **light path diagnostics**, **hot-swap redundant SSDs** and **HDDs**, **power supplies** and **blower modules** with **Calibrated Vectored Cooling™**; **IPMI 2.0** support, including **highly secure remote power control**; **text-console redirect over LAN**, next-generation BIOS (UEFI), an **Advanced Management Module** (upgradeable with a redundant AMM), **IBM Systems Director** management software including **IBM Systems Director Active Energy Manager™**, **Remote Deployment Manager** and **IBM ServerGuide™** help maintain system availability with increased uptime.

If you need scalable, highly manageable, high-performance virtualization in a space- or power-constrained environment, the HX5 is the ideal system.

Selling Features

Price/Performance

- There is an HX5 model to fit all budgets. The HX5 offers a choice of high-performance **4-, 6-, and 8-core Xeon** processors with dual integrated memory controllers, clock rates of **1.86GHz to 2.26GHz**, and **12MB, 18MB or 24MB** of integrated Level 3 cache. Xeon 7500 series processors offer up to **3X** better performance² than the previous-generation 7400 series processors and up to **10X** better performance than the single-core processors of a few years ago that you may still be using.
- Up to **128GB** of **registered DDR3 ECC** memory per blade operates at **800MHz to 1066MHz** (depending on the system configuration), for high performance and wide memory bandwidth.
- By scaling to **2 blade servers**, the HX5 can support up to **256GB** of memory—enough for even the most demanding virtualization, database, or ERP needs. If even that's not enough memory, using **one** blade server and **one memory expansion blades**, the HX5 can scale to **320GB**.
- **Embedded virtualization** (optional on all models) offers extremely high performance, enhanced security, and a zero-gigabyte HDD footprint. (In other words, no mechanical HDD to fail.)

¹ Actual data transfer speed will vary and is often less than the maximum possible. Gigabit Ethernet transfer speed requires support on both system and server, and appropriate network infrastructure.

² Based on Intel measurements.

- **Solid state drives** (SSDs) use only 2W of energy per drive, vs. **9-10W** for 2.5-inch HDDs. This is as much as **80%** less energy than a HDD would use (with a corresponding reduction in heat output).
- The **extremely high degree of integration** in the various BladeCenter chassis reduces the need for server components, replacing numerous fans, KVM and Ethernet cables, power supplies, external switches and other components with fewer *shared* hot-swap/redundant components in the BladeCenter chassis itself. This integration also can greatly **reduce** the amount of **power consumed** and **heat produced**, relative to an equivalent number of 1U servers. This can significantly reduce a data center power bill. The **reduced data center footprint** can also save on infrastructure cost.
- The midplanes used in all chassis provide high-speed **blade-to-blade, blade-to-switch-module and module-to-module communications** internally as well as externally. The midplanes used in the BladeCenter H and BladeCenter HT chassis provide **four 10Gb** data channels to each blade, and supports high-speed switch modules, including **4X InfiniBand** and **10Gb Ethernet**.
- The various BladeCenter chassis use **ultrahigh efficiency power supplies**. Most industry-standard servers use power supplies that are between **70-75% efficient** at converting power from AC wall current to the DC power used inside servers. BladeCenter power modules are up to **92% efficient**. This helps save even more money, as more of the power input you are paying for is used for processing, rather than released into the data center as waste heat that requires even more energy to cool.
- BladeCenter design also **reduces the number of parts required** to run the system. Sharing fans, systems management, and optical media means fewer parts to buy and maintain, and fewer items that can fail and bring the overall solution down.

Flexibility

The **HX5** has the ability to grow with your application requirements, thanks to:

- Up to **two multi-core** Xeon processors (up to **16** cores) per HX5 blade server; up to **2** blades (**4** processors, **32** cores) per server.
- The ability to grow from a single-wide **2-socket** server blade to a double-wide **4-socket** server configuration, optimized for **compute-intensive workloads** with up to 4 processors/32 cores, 256GB of memory, 4 PCIe cards, 16 I/O ports, and 4 SSDs), or for **memory-intensive workloads** with up to one server blade and one MAX5 memory expansion blade (2 processors/16 cores, 320GB of memory, 4 PCIe cards, 16 I/O ports, and 4 SSDs).
- A choice of *processor* speeds (**1.86** or **2.0GHz**), and shared L3 cache sizes (**12MB**, **18MB**, or **24MB**).
- Up to **128GB** of system memory in **16** DIMM slots. Using **two** blade servers the HX5 can support up to **256GB**; or up to 320GB using **one HX5** and **one MAX5** memory expansion blade.
- Up to two internal **hot-swap 1.8-inch** solid state drives, and access to terabytes of external storage via the BladeCenter S chassis or **IBM System Storage™** SAN and NAS storage devices. SSDs consume only 10-20% of the energy required by 2.5-inch HDDs.
- **Two** Gigabit Ethernet ports standard, plus optional **2-port** or **4-port expansion cards** or a **BladeCenter PCI Express I/O Expansion Unit 3**. One HX5 model includes a **10Gb** expansion card (optional in the other models).

In addition, the various **BladeCenter chassis** offer a high degree of flexibility:

- They support configurations that are **30mm** or **60mm** wide, with a variety of I/O options, depending on need.
- When installed in a BladeCenter H high-speed switch bay, the optional **Multi-Switch Interconnect Module** doubles the number of **Gigabit Ethernet** and **Fibre Channel** connections to every blade in the chassis (up to **8** or **12** ports, depending on the blade server).
- Xeon processor-based **HX5** blades can be used in the same chassis as Intel processor-based **HC10, HS12, HS20, HS21, HS21 XM, HS22, HS22V, and HS40** blades; AMD Opteron processor-based **LS20, LS21, LS22, LS41 and LS42** blades; IBM PowerPC® processor-based **JS20, JS21, and JS22 blades**; POWER6® processor-based **JS12, JS23, and JS43 Express** blades; POWER7® processor-based PS700, PS701, and PS702 blades; Cell Broadband Engine™ processor-based **QS21**; and IBM PowerXCell™ processor-based **QS22** blades. Depending on the blade servers used, the various BladeCenter chassis support **Microsoft Windows, Linux, Novell Netware, IBM AIX®** and Sun **Solaris 10** operating systems in the same chassis.
- *Most HS/LS/JS/QS blade servers ever released by IBM are supported in every BladeCenter chassis ever released, going back to 2002³. Every switch module released by IBM is equally compatible. (Ask HP and Dell how far back their compatibility goes.)*
- A blade server has access to as many as **10 communication switches/bridges** in a **BladeCenter**

³ Some older chassis may require power module and management module upgrades.

H or **8** in a **BladeCenter HT** chassis. (Up to **4** switches in a BladeCenter E, S, or T chassis.) And the switches can be Ethernet, iSCSI, SAS, InfiniBand, Fibre Channel, or anything else designed and ServerProven® for BladeCenter use. Switches, bridges and interface cards are currently available from such vendors as Brocade, Cisco, Intel, McData, Nortel, QLogic, Cisco Topspin and others, in addition to IBM.

Manageability

- The HX5 blade server includes an **Integrated Management Module (IMM)** to monitor server availability, perform Predictive Failure Analysis, etc., and trigger IBM Systems Director alerts. The IMM performs the functions of both the Baseboard Management Controller (**BMC**) of earlier systems, and the **Remote Supervisor Adapter II**, as well as **remote presence/cKVM**.
 - Each BladeCenter chassis includes an **Advanced Management Module** to provide additional systems management capabilities, including Web-based out-of-band control; virtual floppy and CD-ROM support; latest OS failure screen capture; LDAP and SSL support; and remote redirection of video, text, keyboard and mouse.
 - Integrated industry-standard Unified Extensible Firmware Interface (**UEFI**) next-generation BIOS. New capabilities include:
 - ☐ Human readable event logs – no more beep codes
 - ☐ Complete setup solution by allowing adapter configuration function to be moved into UEFI
 - ☐ Complete out-of-band coverage by the Advance Settings Utility to simplify remote setup
 - Integrated industry-standard **IPMI 2.0** support works with the IMM to alert IBM Systems Director to anomalous environmental factors, such as voltage and thermal conditions. It also supports **highly secure remote power control**.
 - Integrated **Trusted Platform Module (TPM)** 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.)
 - **IBM Systems Director Active Energy Manager**, an IBM-exclusive, is designed to take advantage of new system power management features, by monitoring *actual* power usage and providing power consumption capping features. More accurate power usage data helps with data center construction planning and the sizing of power and cooling needs, as well as allowing you to use available power more efficiently.
 - The HX5 supports **remote presence/concurrent KVM (cKVM)** and **concurrent media (cMedia)** access by multiple administrators at once, via the IMM.
 - IBM Systems Director is included for proactive systems management and works with both the blade's internal IMM and the chassis' management module. It comes with a portfolio of tools, including IBM Systems Director Active Energy Manager, Service and Support Manager, and others. In addition, IBM Systems Director offers extended systems management tools for additional server management and increased availability. When a problem is encountered, IBM Systems Director can issue administrator alerts via e-mail, pager, and other methods.
-

Availability and Serviceability

- BladeCenter chassis are designed for operation with **greatly reduced potential for single points of failure**. Most aspects of operation, from blade servers to communication modules, to management modules, to power and blower modules, are **hot-swappable**. The midplane connections are **redundant** and the other features can be made so, when used in pairs.
 - Fast **PC3-10600 DDR3 ECC** memory offers multiple layers of Active Memory protection, including **Chipkill** error correction, **Memory ProteXion** (redundant bit steering), **memory scrubbing**, **memory rank sparing**, and **memory mirroring** for high availability
 - HX5 blade servers support the use of **Chipkill-enabled ECC** (error checking and correcting) memory. Chipkill memory can be up to **16X** better than standard ECC memory at correcting some types of memory errors. This can help reduce downtime caused by memory errors.
 - **Solid-state drives** offer up to *triple* the availability (MTBF rates) of conventional SAS HDDs. This can lessen the need for redundant drives.
 - **IPMI 2.0** supports highly secure remote system power on/off using data encryption. This allows an administrator to restart a server without having to visit it in person, saving travel time and getting the server back up and running quickly and securely.
 - **Environmentally tuned blower modules** in the chassis adjust to compensate for changing thermal characteristics. At the lower speeds they draw less power and suffer less wear. Equally important in a crowded data center, temperature-controlled blowers produce less ambient noise in the data center than if they were constantly running at full speed.
 - **Text and graphics console redirection** support allows the administrator to remotely view HX5 text and graphic messages over serial or LAN connections.
-

-
- A **standard three-year (parts and labor) limited onsite warranty**⁴ affords you peace of mind and greater potential investment protection.
-

Right, Open, Easy, Green

You need to make IT decisions that will drive business success. You face management challenges and technological complexity such as space constraints, power and cooling limitations, heterogeneous environments and I/O connectivity issues. IBM brings together the widest choice of compatible chassis, blade servers, storage and networking offerings and solution providers in the industry to help you build an open and flexible IT environment. And regardless of the size of your business, you want to be up and running 24/7. With built-in redundancy, innovative power and cooling and the latest I/O and management tools, IBM BladeCenter is easy to own—so you can focus on your business demands and stay ahead of the competition.

The **RIGHT** choice, tailored to fit your diverse needs:

- It's flexible and modular. As needs evolve, a one-size-fits-all solution doesn't work.
 - Meet your needs with BladeCenter: everything from a high-performance data center to a small office with limited IT skills—IBM has you covered
 - Get flexibility with 5 compatible chassis and 5 blade types supporting multiple I/O fabrics, all managed from a common point
- It's robust and reliable, providing redundancy throughout and the information you need to keep your business up and running.
 - Provide redundancy for no single point of failure with IBM BladeCenter
 - Preserve application uptime with IBM Predictive Failure Analysis and light path diagnostics
 - Make decisions based on accurate data for quick problem diagnosis with First Failure Data Capture

OPEN and innovative, for a flexible business foundation:

- It's comprehensive, providing broad, fast, and reliable networking and storage I/O with BladeCenter Open Fabric.
 - Match your data center needs and the appropriate interconnect using a common management point, and 5 I/O fabrics to choose from
 - Extract the most from your third-party management solutions by utilizing the BladeCenter Open Fabric Manager
- It's collaborative, enabling you to harness the power of the industry to deliver innovation that matters.
 - Get flexibility from a myriad of solutions created by Blade.org members and industry leaders that have downloaded our open specification

EASY to deploy, integrate and manage:

- It enables efficient integrated management, which allows you to minimize costs with the tools you need for effective management.
 - Automate OS installation and BIOS updates remotely with IBM Systems Director tools
 - Administer your blades at the chassis or rack level with the Advanced Management Module
 - Plug into your enterprise management software
- It enable deployment simplicity without tradeoffs by speeding the deployment of new hardware in minutes rather than days, using BladeCenter Open Fabric Manager
 - Get significantly faster deployment of servers and I/O than from rack solutions
 - Reduce costly downtime with integrated failover capability
 - Manage from a single point of control via the Advanced Management Module
 - Use with virtually all IBM switches, blades and chassis

GREEN today for a better tomorrow:

- It offers control via powerful tools that help you optimize your data center infrastructure so you can be responsive.
 - Understand your power requirements with IBM Power Configurator
 - Monitor, control and virtualize your power with IBM Systems Director Active Energy Manager
 - Reduce data center hot spots with the IBM Rear Door Heat eXchanger
 - Optimize and future-proof your data center with IBM Data Center Energy Efficiency services
 - Our eco-friendly servers and services can help you be environmentally responsible.
-

⁴ For terms and conditions or copies of the IBM Statement of Limited Warranty, call 800-772-2227 in the U.S. In Canada call 800-426-2255. Telephone support may be subject to additional charges. For warranties including onsite labor, a technician is sent after IBM attempts to resolve the problem remotely. International warranty service is available in any country in which this product is sold.

- Become more energy efficient with IBM expertise

Key Features

Multicore Intel Xeon Processors

The HX5 ships with 1 or 2 high-performance Intel **Xeon 7500** or **6500 Series** processors per blade. By connecting 2 blades together, the HX5 servers also go a step further by allowing you to increase the number of sockets from 2 to as many as 4. The choice of processors includes:

- **130W 8-core** Xeon processor model **X7560** at 2.26GHz, with 64-bit extensions, *low power draw per core (16.25W)*, **6.4 GTps** (gigatransfers per second) QPI speed, **978MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **130W 8-core** Xeon processor model **X6550*** at 2.0GHz, with 64-bit extensions, *low power draw per core (16.25W)*, **6.4 GTps** QPI speed, **978MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **95W 8-core** Xeon processor model **L7555** at 1.86GHz, with 64-bit extensions, extremely *low power draw per core (11.9W)*, **5.86 GTps** QPI speed, **978MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **105W 6-core** Xeon processor model **E7540** at 2.0GHz, with 64-bit extensions, *low power draw per core (17.5W)*, **6.4 GTps** QPI speed, **978MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **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, **978MHz** memory access, dual integrated memory controllers, **12MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **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, **18MB** of shared **L3** cache, and Intel Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis
- **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; supported in BladeCenter H, HT, and S chassis
- **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 (available via CTO), and Intel Hyper-Threading technology; supported in BladeCenter H, HT, and S chassis

* This processor is limited to single-chassis configurations.

Note: Because of the integrated memory controllers the former front-side bus (FSB) no longer exists.

With the Xeon 7500 series processors, Intel has diverged from its traditional Symmetric Multiprocessing (SMP) architecture to a Non-Uniform Memory Access (NUMA) architecture. The Xeon 7500 processors are connected through three serial coherency links called QuickPath Interconnect (QPI). QPI is capable of up to **6.4GTps** (gigatransfers per second), depending on the processor model.

The **4-core Xeon** processors contain **4 complete processor cores** and **8 threads**. Each processor contains one L3 cache, shared by all the cores. The cores appear to software as separate physical processors. Four-core processors can offer more than double the performance of a same-speed 2-core Xeon processor (depending on workload). Similarly, **8-core** Xeon processors contain **8 processor cores** and **16 threads**. They can offer more than double the performance of four-core Xeon processors (again, depending on workload). Likewise, the **six-core** processors contain **six** processor cores.

Each processor includes two integrated memory controllers, to reduce memory bottlenecks and improve performance.

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

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

DDR3 Registered Memory with Active Memory Protection

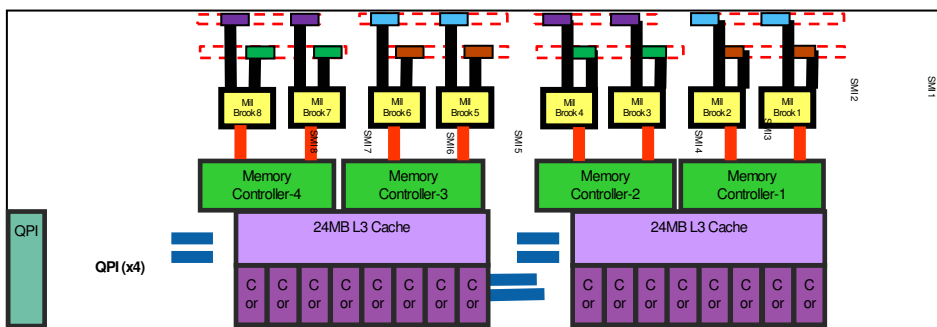
The HX5 uses registered double data rate III (DDR3) VLP (very-low-profile) DIMMs and provides Active Memory features, including advanced **Chipkill** memory protection, for **up to 16X** better error correction than standard ECC memory. In addition to offering **triple the memory bandwidth** of DDR2 or fully-buffered memory, DDR3 memory also uses less energy. DDR2 memory already offered up to **37%** lower energy use than fully buffered memory. Now, a generation later, DDR3 memory is even more efficient, using **10-15% less energy** than DDR2 memory.

The HX5 supports up to **128GB** of memory in **16** DIMM slots. Redesign in the architecture of the Xeon 7500 series processors bring radical changes in the way memory works in these servers. For example, the 7500 series processors **integrate 2 memory controllers inside each processor**, resulting in four memory controllers in a two-socket system. Each memory controller has four memory channels, one per pair of DIMMs. Depending on the type of memory, population of memory, and processor model, the memory may be clocked at **978MHz** or **800MHz**. Adding a *second* HX5 server to the first doubles the memory capacity to **256GB** in **32** DIMM slots.

Redesign in the architecture of the x7500 series processors bring radical changes in the way memory works in these servers. For example, the Xeon 7500 series processors **integrate two memory controllers inside each processor**, resulting in **four** memory controllers in a two-socket system. Each processor has four memory channels.

Note: If only one processor is installed, only **eight** DIMM slots (up to 64GB) can be used. There are two ways to expand memory beyond 8 DIMMs. Adding a second processor not only doubles the amount of memory available for use, but also doubles the number of memory controllers, thus doubling the system memory bandwidth. If you add a second processor, but no additional memory for the second processor, the second processor has to access the memory from the first processor “remotely,” resulting in longer latencies and lower performance. The latency to access remote memory is almost **75% higher** than local memory access. So, the goal should be to always populate both processors with memory.

Alternatively, you can add a **MAX5** memory expansion blade, containing **24** more DIMM slots with up to **192GB** of memory. This offers up to **32** DIMMs and **256GB** of memory with only a **single processor**. Adding the second processor allows the use of **8** more DIMM slots inside the HX5 and a grand total of **40** DIMMs and **320GB** of memory in a double-wide (60mm) blade server. Using a MAX5 unit also boosts memory performance to **1066MHz** for the server.



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:

- Always populate processors with equal amounts of memory to enable a balanced NUMA system
- Always populate both memory channels on each processor with equal memory capacity
- Ensure an even number of ranks are populated per channel

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

- Xeon 7500 Series 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.

With **memory rank sparing**, two unused 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.

Memory mirroring works much like disk mirroring. The total memory is divided into two channels. Data is *written concurrently to both channels*. If a DIMM fails in one of the DIMMs in the primary channel, it is instantly disabled and the mirrored (backup) memory in the other channel becomes active (primary) until the failing DIMM is replaced. With mirroring enabled, one-half of total memory is available for use. (**Note:** Due to the double writes to memory, performance is affected.) Mirroring is handled at the hardware level; no operating system support is required.

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 is available in **2GB, 4GB, 8GB, or 16GB** DIMMs.

High-Performance Scalability

Because of eX5 design, an HX5 server has the incredible flexibility of scaling from **2** sockets to **4** sockets with ease, doubling the memory and storage capacity and I/O slots in the process. A 2-node configuration is achieved simply by connecting **2** 30mm blades together via the Quick Path Interconnect (QPI) ports, for **4** sockets and **32** DIMMs (**256GB**).

Alternatively, you can scale via **MAX5**, adding memory capacity to increase performance in an I/O-intensive environment. This configuration connects **1** HX5 blade to **1** MAX5 blade, for a total of up to **40** DIMMs (**320GB**). **Note:** Selected models of HX5 come standard with a MAX5 blade.

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

Integrated Virtualization

HX5 offers a separately orderable **IBM 2GB USB Flash Drive** option preloaded with **VMware vSphere Hypervisor** (previously ESXi), an embedded version of VMware ESX Server. Microsoft **Hyper-V** and **Xen** are also supported, installed by the user on an SSD.

ESXi is fully contained on the flash drive, without requiring any disk space. 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. ESXi includes all the performance, scalability and compatibility features of ESX Server, including full **VMFS** support across FC SAN, iSCSI SAN, and NAS, and **4-way VSMP**. Rather than management through a Service Console based on a Linux operating system, ESXi relies on aggregate management tools, including VirtualCenter, the Remote Command Line interface and the introduction of CIM for standards-based and agentless hardware monitoring. Licensing works the same as for "standard" ESX Server.

Disk Controllers

HX5 models support an optional **LSI 1064E SAS/SATA** controller (required for SSD support). This controller supports up to **two** internal **solid-state drives** and offers **hardware RAID-0/1/1E** support.

High Performance/High-Availability Internal Storage

The HX5 supports up to **two fixed 1.8-inch solid-state** drives (SSDs) internally:

- **SATA-based solid-state — 50GB capacity (100GB 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 availability of mechanical disk drives
- No moving parts to fail
- Enterprise wear-leveling to extend life even further

❑ **Full OS Support**

- Supports all ServerProven OSes

Due to the statistically higher failure rates for traditional spinning media, IBM recommends the use of the solid state drives (SSDs) as an alternative. They store data on flash memory chips, rather than on magnetic media. Like HDDs, SSDs can be used as boot media and for random access storage. However, SSDs offer higher thresholds of shock and vibration, and a higher operating temperature range (between 0 and 70 degrees C). This yields a failure rate only 1/3 that of HDDs (approximately 3,000,000 hours MTBF vs. 1,000,000 hours). In addition, the **IBM 50GB 1.8" Solid State Drive** requires only **2W**. This reduces the storage power requirement and heat output by as much as **80%**, compared to 2.5-inch HDDs. If used as a boot device, no special device drivers are required.

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**)

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 **500⁵** 3.5-inch 600GB 15K SAS HDDs, with **99% lower latency** (30 µs), and **7.7x** the bandwidth of a HDD. On a performance-per-watt basis, these adapters outperform HDDs by up to **445x⁶**.

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.

Note: The High IOPS Adapters require a BPE4 expansion unit.

Additional storage is available via the BladeCenter S chassis.

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

External Tape Storage

The HX5 supports various external rack-mounted SAN-attached tape drives. Supported tape technologies include:

- IBM System Storage TS2230 Express Tape Drive
- IBM System Storage TS2240 Tape Drive Model S4E
- IBM System Storage TS2900 Tape Autoloader
- IBM System Storage TS3100 Tape Library
- IBM System Storage TS3200 Tape Library
- IBM System Storage TS3310 Modular Tape Library
- IBM System Storage TS3500 Tape Library



⁵ According to internal IBM testing, the typical **maximum IOPS** of a 600GB 15K 3.5-inch hot-swap SAS HDD is **~400** (4K random reads), while the typical IOPS of a 640GB High IOPS SSD PCIe adapter is **~200,000**, or **~500x**. The typical **bandwidth** of the same HDD = **~195MBps** (64K sequential reads), while the typical bandwidth of a 640GB High IOPS MLC adapter is **~1.5GBps**, or **~7.7x**.

⁶ 97,014 IOPS / 9W = **10,779 IOPS** per watt (160GB/320GB adapters). 196,000 IOPS / 12W = **~16,333 IOPS** per watt (640GB adapters). 400 IOPS / 16.5W = **~24 IOPS** per watt (600GB 15K 3.5-inch hot-swap SAS HDD).



Light Path Diagnostics

Light path diagnostics enables a technician to quickly identify and locate a failed or failing system component, such as a specific blower module or memory DIMM. This enables quick replacement of the component, which helps increase server uptime and lower servicing costs.

The front of each blade server—and the chassis itself—has an LED indicator light to show possible component failures. This lets the servicer identify the failing component without the need to or remove the blade server from the chassis. The light path diagnostics panel tells the servicer which component of the affected server 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 both processors. This allows 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.)



Gigabit Ethernet Controller

The HX5 includes a **dual-port** integrated **Broadcom BCM5709S** Gigabit Ethernet controller for up to 10X higher maximum throughput than a 10/100 Ethernet controller. The controller offers **TOE** (TCP Offload Engine) support, as well as **failover** and **load balancing** for better throughput and system availability. It also supports highly secure remote power management using **IPMI 2.0**, plus Wake on LAN® and PXE (Preboot Execution Environment) Flash interface.

If 2 ports aren't enough, optional **2-port** or **4-port Ethernet** expansion cards can be used for additional ports. For example, the **CIOv** slot can hold a **2-port** card, the **CFFh** slot supports a **4-port** card, and a **bridge module** can add another **2 ports**, for a total of **10** Gigabit Ethernet ports per HX5 blade.

High-Performance Adapter Slots

The HX5 blade server includes **two PCIe** adapter slots. They support **CFF** (compact form factor) cards: **one** standard-speed **CIOv** and **one** high-speed **CFFh** card. **Note:** The SSD expansion card is installed in the upper I/O slot. When SSDs are installed, both I/O slots are still available for use.

Adding a second HX5 blade for a 2-wide 4-socket server doubles the number of PCIe card slots.

The **BladeCenter PCI Express I/O Expansion Unit 3 (BPE3)** adds **2 standard full-height/full-length x16** physical/x8 electrical (**4GBps**) **PCIe Gen 1** expansion card slots, supporting adapters of up to **25W** apiece to an HS22. One BPE3 can be connected per HS22 blade, for a total of **3** available slots (**1** in the blade and **2** in the expansion unit). **Note:** The BPE3 reserves the high-speed CFFh expansion connector in the HS22, leaving only the CIOv slot available.

Similarly, the optional **BladeCenter PCI Express Gen 2 Expansion Blade (BPE4)** adds **1** standard **full-height/full length** and **1** standard **full-height/half-length x16** physical/x8 mechanical (**8GBps**) PCIe Gen 2 expansion card slot per HS22 blade. These slots support two industry-standard PCIe adapters, up to **75W** per adapter. The **BladeCenter PCI Express Gen 2 Expansion Blade** offers a unique stacking feature that allows clients to stack up to **4** expansion blades per HS22 blade, offering up to an additional 8 PCIe slots per HS22 blade. **Note:** Unlike the BPE3, the BPE4 does *not* reserve the high-speed CFFh expansion connector in the HS22, leaving *both* slots available in the server.

If I/O slots are a greater need than processors or memory, attaching multiple I/O expansion units to one blade server is much more cost-effective than installing multiple blade servers for the same number of adapter slots.

Adapters can be used to add fabrics to BladeCenter switch modules, including 10Gb Ethernet, additional Gigabit Ethernet controllers, CNA, Fibre Channel, InfiniBand, SAS, etc.

BladeCenter Chassis

IBM's blade architecture offers **five choices** of compatible and interoperable chassis in which to use various blade servers. Each chassis serves different customer needs. The **BladeCenter S** is a small, entry-level chassis designed for office environments. The original chassis (refreshed with the latest Advanced Management Modules and power supply modules) offers maximum density, great flexibility and a wide variety of expansion options at an entry-level price. The next-generation **BladeCenter H** chassis offers all of BladeCenter's capabilities, and adds high-performance features, including 10Gb fabric support. If you need a **ruggedized** chassis (for example, government/military or telecom), **BladeCenter T** offers special features optimized for those environments. The next-generation **BladeCenter HT** is a high-performance **ruggedized** telecommunications platform, also supporting 10Gb fabrics. *HX5 is supported in the BladeCenter*



H, BladeCenter HT and BladeCenter S chassis.

There is a high degree of interchangeability and compatibility of features among the chassis, and all of these chassis can be installed in a rack along with other rack-optimized equipment.

See the separate *IBM BladeCenter Chassis Product Guide* for details.

Advanced Systems Management Capabilities

Each BladeCenter chassis offers a high level of systems management capabilities that are well-suited to remote locations as well as to stand-alone environments. Features include the Advanced Management Module (AMM), Integrated Management Module (IMM), IBM ToolsCenter, Automatic Server Restart, Systems Director Active Energy Manager, Wake on LAN support, PXE 2.0 support, text and graphics console redirect, Predictive Failure Analysis, Embedded Dynamic System Analysis, and IBM Systems Director.

The chassis-based **AMM**, in combination with the HX5 blade server's **IMM**, provides industry-standard **Intelligent Platform Management Interface (IPMI) 2.0**-compliant systems management. The AMM 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
- 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 functions in the blade's IMM

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

Other systems management features offered for the combination of blade server and chassis include:

- Predictive Failure Analysis for system processors, memory and HDDs, as well as chassis switch modules, blower modules and power modules
- Web-based out-of-band control
- Latest OS failure screen capture
- Remote virtual media
- High-speed remote redirection of PCI video, keyboard and mouse
- SSL (Secure Socket Layer) and LDAP (Lightweight Directory Access Protocol) support

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.

In order to put control of processor power-saving features at the fingertips of administrators, IBM developed **IBM Systems Director Active Energy Manager**. Active Energy Manager is designed to take advantage of new processor features, such as balancing the performance of the system according to available power input. It provides the ability to plan and predict power consumption based on your BladeCenter hardware configuration. It also helps you 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 at the blade level, then adding up the power draw for each blade and tracking that usage. In failure mode, Active Energy Manager (through the BladeCenter Management Module) might request that certain blades in each domain throttle down to reduce power consumption.

Automatic Server Restart (ASR) helps reduce downtime by restarting the server automatically in the event of a system lockup. ASR technology is a combination of hardware circuitry tied into the server's system reset function and a device driver. As long as the server continues running, the ASR watchdog timer will keep being reset, but if the operating system crashes or the hardware freezes somehow the ASR software will be unable to reset the hardware timer. If the timer is not reset within five minutes, 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*.

Text and Graphics Console Redirect support allows the administrator to remotely view HX5 text and graphics messages over serial or LAN.

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 allows software such as the **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.

Predictive Failure Analysis (PFA) enables the AMM and the IMM to detect impending failure of supported components (processors; memory; expansion cards; switch, blower and power supplies; 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.

Embedded Dynamic System Analysis (DSA) feature collects and analyzes system information to aid in diagnosing problems. The diagnostic programs collect the following information:

- System configuration
- Network interfaces and settings
- Installed hardware
- Service processor status and configuration
- Vital product data, firmware, and BIOS configuration
- Hard disk drive health
- RAID controller configuration
- Event logs for ServeRAID controllers and service processors

The diagnostic programs create a merged log that includes events from all collected logs. The information is collected into a file that you can send to IBM service and support. Additionally, you can view the information locally through a generated text report file, or you can copy the log to removable media and view the log from a Web browser.

IBM Systems Director software for advanced workgroup management is included with the server. IBM Systems Director comes with a portfolio of tools, including *Active Energy Manager*, *Service and Support Manager*, and other tools. *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.



Extensive System Support Features

The IBM services and technical support portfolio provides world-class, consistent, high-quality service and support. From the start, IBM programs make it easier for you to plan for, configure and purchase BladeCenter servers, get them running and keep them running long-term. These features include IBM ServerProven, 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 blade servers and chassis and are officially supported to work together. It is updated frequently to keep the latest compatibility information 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 servers. ServerGuide goes beyond mere hardware configuration by assisting with the automated installation of the Microsoft Windows Server® 2008 operating system, 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 to 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 service⁷ 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 **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 both 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 let you take your servers to a higher level

You rely on System x and blade options to supply a comprehensive solution for your business needs. Options help 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 processor provides high clock rates, 64-bit extensions, a large cache, Hyper Threading Technology, Turbo Boost Technology, energy-saving features, and

⁷ For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

advanced features for availability and manageability. Large cache size and multiple cores, combined with an **800MHz** or **978MHz** memory access speed, reduce memory latency and facilitate 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, as well as the number, capacity, and arrangement of the DIMMs.) 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 BladeCenter server is one of the most effective ways to increase application performance. For best performance in a server with a 4-core processor, there should be twice as much memory available as for a 2-core processor. With 8 cores, there should be twice as much memory as for a 4-core processor. **Note:** Fewer large-capacity DIMMs may provide better performance and use less energy than more lower-capacity DIMMs. On the other hand, more lower-capacity DIMMs can be less expensive to buy than fewer high-capacity DIMMs.

DIMM slots — Additional slots are available either by adding a second HX5 blade to the first, to form a **32-DIMM-slot 2-node server**, or by adding a MAX5 memory expansion unit to the HX5 blade for **40** total DIMMs slots.

I/O slots — Additional slots are available by adding a second HX5 blade to the first, to form a 2-node server.

I/O Expansion Units — A choice of BladeCenter **PCI Express Expansion Unit 3 (BPE3)** or BladeCenter **PCI Express Gen 2 Expansion Blade (BPE4)** offers from **3** to **10** total adapter slots per HS22 server, to support additional fabrics (including Ethernet, 10Gb Ethernet, Fibre Channel, SAS, and Converged Network Adapters), and high-performance graphics processing units (GPUs).

Solid State Drives — IBM offers a **50.0GB 1.8" solid-state drive** as a higher reliability, lower-energy-use 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. Additional direct-attach storage is available via the BladeCenter S chassis.

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, with **160GB** or **320GB** capacities.

ServeRAID Controllers — System x servers using ServeRAID technology allow organizations 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.

External Storage — The IBM **System Storage DS3000**, **DS4000**, **DS5000**, and **DS8000** series, as well as the **System Storage N3000**, **N5000**, **N6000**, and **N7000** series, comprise a powerful and broad shared storage family with integrated management software designed to meet midrange and enterprise needs. For lower-end needs, IBM offers the System Storage **EXP810**, and **EXP3000** storage expansion units. Additionally, external LAN-attached tape storage is available.

Communication Modules — The various BladeCenter chassis support integrated communication and I/O switches and/or bridges for **Gigabit** and **10Gb Ethernet**, **Myricom**, **4Gb** and **8Gb Fibre Channel**, **1X** and **4X InfiniBand**, **iSCSI**, **3Gb SAS**, and others. Expansion adapters for individual HX5 blades are available to interface with these modules. The BladeCenter H and BladeCenter HT also support the **Multi-Switch Interconnect Module**, which doubles the number of Gigabit Ethernet and Fibre Channel connections available to all blades in the chassis.


Rear Door Heat eXchanger — The unit attaches to the back of an IBM S2 42U Enterprise Rack. It is capable of removing up to 50,000 BTUs (14KVa) from the data center using water lines under the raised floor. The door swings open for servicing.

Redundant features — Optional hot-swap drives, power supply modules, blower modules, management modules, switches and bridges provide redundancy for the various BladeCenter chassis.


Optimized for virtualization and database applications with maximum memory and compute capacity in a blade

HX5 Images

Front View



Interior View



Optimized for virtualization and database applications with maximum memory and compute capacity in a blade

BladeCenter HX5 Specifications			
Machine type	7872-42x, 61x/63x/64x/65x/66x/67x/68x/6Dx, 82x/83x/84x/86x, E6x/E8x		
Form factor (standard / maximum)	30mm / 120mm (2 30mm HX5 blades or 1 30mm HX5 blade + 1 30mm MAX5 memory expansion blade)		
Processor type	Eight-core Xeon (X75xx/L75xx/X65xx) 1.86GHz L7555 (82x), 2.0MHz, X6550 (83x, E8x), 2.26GHz X7560 (84x/86x)	Six-core Xeon (E75xx/E65xx) 1.86GHz E7530 (61x), 2.0GHz E6540 (63x/67x/68x, 6Dx), 2.0GHz E7540 (64x/65x/66x, E6x)	Four-core Xeon (E75xx/E65xx) 1.73GHz E6510 (CTO), 1.86GHz E7520 (42x)
Processor cores per 2-socket blade (30mm)	16 (82x/83x/84x/86x, E8x)	12 (61x/63x/64x/65x/ 66x/67x/68x/6Dx, E6Y)	8 (42x)
Processor cores per 4-socket blade server (60mm)	32 (82x/83x/84x/86x, E8x)	24 (61x/63x/64x/65x/ 66x/67x/68x/6Dx, E6Y)	16 (42x)
Maximum processor power draw	130W (83x/84x/86x, E8x)	105W (61x/63x/64x/65x/ 66x/67x/68x/6Dx, E6Y)	95W (42x, 82x)
QuickPath Interconnect (QPI) speed (gigatransfers per second)	6.4 GTps (64x/65x, 83x/84x/86x, E6Y, E8x)	5.86 GTps (61x/63x/65x/ 66x/67x/68x/6Dx, 82x)	4.8 GTps (42x)
# of processors standard / maximum per blade and per 2-node system	1 / 2 ; up to 4 total (2-node configuration)— 42x, 61x/64x/65x, 82x/86x		2 / 2 ; up to 4 total (2-node configuration)— 63x/66x/67x/68x/6Dx, 83x/84x, E6x/E8x
Internal L3 cache	24MB (82x/83x/84x/86x, E8x)	18MB (42x, 63x/64x/65x/66x/67x/68x/6Dx, E6x)	12MB (61x, E6510)
Chipset	Intel 7500 / IBM eX5		
Scalability ports	4 QPI Ports per chassis 3 EXA ports per MAX5		
# of DIMM sockets total / available	16 / 14		16 / 0
Standard / maximum memory^a per blade	160GB (40 x 4GB) / 128GB (67x/68x)	112GB (28 x 4GB) / 128GB (E8x)	64GB (16 x 4GB) / 128GB (E6x)
	32GB (8 x 4GB) / 128GB (66x)	16GB (4 x 4GB) / 128GB (63x/6Dx, 83x/84x)	8GB (2 x 4GB) / 128GB (42x, 61x/64x/65x, 82x/86x)
Maximum memory, HX5 blade plus 1 MAX5 expansion blade	320GB (40 x 8GB)		
Maximum memory, 2 HX5 blades	256GB (32 x 8GB)		
MAX5 expansion blades included	Standard—63x/6Dx/67x/68x, 83x/84x, E8x		Optional—All other models
Standard memory type	Registered PC3-10600 DDR3 ECC (1066MHz)		
Maximum memory speed	978MHz—82x/83x/84x/ 86x, E8x (1066MHz with MAX5)	978MHz—61x/63x/64x/65x/ 66x/67x/68x/6Dx, E6x	800MHz—42x, E6510 (via CTO)
Memory interleaving	Yes		
DIMM capacities supported	2GB, 4GB, 8GB VLP		
Chipkill protection supported	Yes		
Memory ProteXion supported	Yes		
Memory mirroring supported	Yes		

^a Maximum memory and disk capacity may require the replacement of standard components with the largest supported component available.

Optimized for virtualization and database applications with maximum memory and compute capacity in a blade

BladeCenter HX5 Specifications			
Memory rank sparing supported	Yes		
Memory scrubbing supported	Yes		
# of 1.8-inch drive bays total / available	2 / 2 SSD (per server blade)	4 / 4 SSD (per 4-socket server configuration)	
# of 2.5-inch drive bays total / available	None		
# of 3.5-inch drive bays total / available	None		
Maximum internal SSD capacity	100GB (2 x 50GB) fixed	200GB (4 x 50GB) fixed	
SSD capacities supported	50.0GB		
# of SSDs standard	None		
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		
High IOPS Adapters standard	2 x 320GB High IOPS MS Class SSD PCIe Adapter—66x	None—all other models	
Modular Flash Drives supported	1		
Internal USB Flash Drive supported	1 standard (preloaded with VMware vSphere 4.1 hypervisor)—68x	1 standard (preloaded with RHEV-H hypervisor)—67x	1 optional (preloaded with hypervisor)—all other models
# of optical drives standard	None (one in BladeCenter chassis)		
# of diskette drives standard	None (one standard in BladeCenter H chassis)		
Internal tape drives supported	None (SAN-attached)		
SSD drive interface	SATA		
Integrated disk controller	SATA controller integrated into the processor		
Optional RAID controller	LSI Logic 1064		
RAID levels supported	RAID-0/1/1E (optional)		
External disk drive support	NAS/SAN-attach		
# of PCIe slots	1 CIOv and 1 CFFh per node	2 CIOv and 2 CFFh per 4-socket/2-node server configuration	
# of PCIe expansion slots	Up to 3 total, using 1 BladeCenter PCI Express I/O Expansion Unit 3 (BPE3) blade with one HX5 blade	Up to 10 total using 4 BladeCenter PCI Express Gen 2 Expansion Blade (BPE4) blades with one HX5 blade	
PCIe expansion units standard	BPE4—66x	None—all other models	
# of I/O ports per blade (maximum)	8 (16 in a 4-socket server configuration)		
# of video ports	None (chassis-attached)		
Video controller	Matrox G200eV (in IMM)		
Video memory	16MB SDRAM		
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		
Gigabit Ethernet controllers std	Dual-port Broadcom BCM5709S —TOE-enabled		

Optimized for virtualization and database applications with maximum memory and compute capacity in a blade

BladeCenter HX5 Specifications		
# of Gigabit Ethernet ports	2 (standard per blade); 12 with BladeCenter Virtual Fabric Adapter	
# of RS485 ports	None	
# of serial ports	None (1 direct via BladeCenter H chassis, or Serial over LAN in BladeCenter H)	
# of parallel ports	None	
# of mouse ports	None (1 via chassis)	
# of keyboard ports	None (1 via chassis)	
# of USB ports	1 (plus 2 via chassis)	
Systems management controller	Integrated Management Module	
Diagnostic LEDs (front panel)	Power good, blade location, over temperature, information, general fault	
Predictive Failure Analysis support	Processor, memory, SSDs, PCIe expansion cards	
BIOS type	UEFI 2.1	
Power supply size	Contained in chassis	
# of power supplies standard / maximum	Contained in chassis	
# of fans/blowers standard / maximum	Contained in chassis	
Dimensions (HWD) / weight	9.7" (245mm) H 1.14" (29mm) W 17.6" (446mm) D	12 lbs (maximum) 5.44 kg (maximum)
Operating systems supported	Microsoft Windows Server 2008 and R2 (Standard/Web/Enterprise/Datacenter Editions) 32/64-bit, Windows HPC Server 2008 64-bit, RHEL 6 64-bit—without Xen, RHEL 5.3 64-bit—with or without Xen, SLES 11 64-bit with or without Xen, SLES 10 64-bit with or without Xen, VMware ESX/ESXi/vSphere Hypervisor 4.0	
Extended/long life support	Selected blades are supported for long life. (Contact IBM sales for details.)	
Length of limited warranty	3 years (parts and labor) onsite	

The Bottom Line

The HX5 offers maximum bang for the buck by incorporating industry-leading features in a tiny package:

Price/Performance

- **Multi-core processors** — Up to **two 1.86GHz 4-core**, **1.86 to 2.26GHz 6-core**, or **1.86GHz 4-core** Xeon **7500** Series processors; up to **two 1.73GHz 4-core** Xeon **6500** Series processors (via CTO)
- **Large cache** — **12MB to 24MB** of L3 processor cache (processor-specific)
- **Huge memory capacity** — Up to **128GB** per blade; up to **320GB** (using MAX5)
- **Fast memory access** — **978MHz** or **800MHz** access (depending on the processor model and DIMMs used); up to 1066MHz, using MAX5
- **Embedded virtualization** (Separately orderable option)
- **High-IOPS technology** — Optional **High IOPS** SSD adapters
- **Fast I/O** — 10Gb Ethernet, 4X InfiniBand, iSCSI, 6Gb SAS or 8Gb Fibre Channel connectivity via an optional expansion card
- **Fast communications** — Integrated **dual Gigabit** Ethernet, with the option for two or four additional Gigabit Ethernet ports via expansion cards and expansion unit (total **12** ports per blade); support for **10Gb Ethernet**

Flexibility

- **Single-** and **double-wide** configurations (up to 2 blade servers or 1 blade server and 1 MAX5 memory expansion unit)
- **Flexible memory configurations** — **16** DIMMs per blade; **32** DIMMs with two blades; **40** DIMMs

- with 1 blade and 1 MAX5 expansion unit
- Up to **two** internal **1.8-inch non-hot-swap** SSDs per node
- **Solid-state drives** are an energy-saving, higher-reliability alternative to HDD storage
- Optional IBM **160GB, 320GB, or 640GB High IOPS SSD Adapters** for extremely high IOPS and bandwidth storage (requires BPE4 expansion unit)
- Integrated RAID — **RAID-1** mirrored arrays optional
- **Two available adapter slots** standard :
 - ☐ **Two** CFF (compact form factor) slots: one high-speed CFFh and one standard-speed CIOv
- Additional **optional** adapter slots⁹:
 - ☐ Up to **3** total, using **1** BladeCenter **PCI Express I/O Expansion Unit 3** blade
 - ☐ Up to **10** total, using **4** BladeCenter **PCI Express Gen 2 Expansion Blades** (BPE4) blades

Manageability and Availability

- **IBM Systems Director** systems management software, including:
 - ☐ IBM Systems Director **Active Energy Manager**
 - ☐ IBM **Service and Support Manager**
- **Integrated Management Module:**
 - ☐ **IPMI 2.0** compliance, including highly secure remote power control
 - ☐ **cKVM**
 - ☐ **Advanced management capabilities**
- Interface to one or two **Advanced Management Modules** in the chassis for advanced systems management capability
 - ☐ Supports **LDAP** and **SSL** industry standards
 - ☐ **Text and graphics console redirection** systems management
 - ☐ **Serial over LAN**
- **Active Memory protection:**
 - ☐ Advanced **Chipkill** ECC memory correction
 - ☐ **Memory ProteXion** support
 - ☐ Memory **rank sparing**
 - ☐ Memory **mirroring**
 - ☐ Memory **scrubbing**
- **Solid-state drive** as a high-reliability alternative to internal HDD storage (with up to three times the MTBF of spinning disk drives)
- Numerous **hot-swap/redundant capabilities** provided via the chassis

⁹ The PCIe expansion units each require an adjacent blade slot.

Optimized for virtualization and database applications with maximum memory and compute capacity in a blade

Server Comparison Chart

The following table shows the suggested uses for the respective IBM BladeCenter blade servers, including comparisons of the uses for which each server is best suited:

		Requirements										BladeCenter Servers											

For More Information

IBM BladeCenter Servers and Options
IBM Systems Director Service and Support Manager
IBM System x and BladeCenter Power Configurator
IBM Standalone Solutions Configuration Tool
IBM Configuration and Options Guide
IBM ServerProven Program
Technical Support
Other Technical Support Resources

<http://ibm.com/systems/bladecenter>
<http://ibm.com/support/electronic>
<http://ibm.com/systems/bladecenter/resources/powerconfig.html>
<http://ibm.com/systems/x/hardware/configtools.html>
<http://ibm.com/systems/x/hardware/configtools.html>
<http://ibm.com/systems/info/x86servers/serverproven/compat/us>
<http://ibm.com/server/support>
<http://ibm.com/systems/support>

Legal Information

© IBM Corporation 2010

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

Produced in the USA
October 2010
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 Memory, AIX, BladeCenter, Calibrated Vectored Cooling, Chipkill, DB2, POWER6, POWER7, PowerPC, PowerXCell, ServerGuide, ServeRAID, System Storage, System x, Systems Director Active Energy Manager, Wake on LAN, and X-Architecture 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>.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

InfiniBand is a trademark of the InfiniBand Trade Association.

Intel and 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, Windows, the Windows logo, Windows Server, and SQL Server 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.

Some machines are designed with a power management capability to provide customers with the maximum uptime possible for their systems. In extended thermal conditions, rather than shutdown completely, or fail, these machines automatically reduce the processor frequency to maintain acceptable thermal levels.

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.

BLO03039-USEN-04