



Innovation Sets IBM System x 1P and 2P Servers Apart from the Competition

*By Mark T. Chapman
IBM Systems and Technology Group*

Executive Overview

In tough economic times like these, when you need more computing capacity it's easy to say, "Order some more commodity low-end servers. They're cheap, so we can save money. Besides, they're all the same anyway." But is this really true? Are uni- and dual-processor servers all the same?

In fact, they're not. There are many differences between cheap cookie-cutter bare-bones servers and servers loaded with features designed to run reliably and save you money. What's most important isn't the entry hardware cost, but the *long-term cost of operating that hardware*. You should ask yourself, "What are they leaving out of those cheap servers to keep the purchase price down?"

It may be systems management hardware and software that can save you money on administration costs. It may be redundant hot-swap and simple-swap hardware that can increase availability and reduce service time. It may be energy-efficient components that can save you money on power and cooling. It may be toolless components that can save you time during installation. It may be extensive reliability and compatibility testing that helps maximize system uptime. There are many design features that add slightly to the initial hardware cost but can save you a lot more money in the long run. And isn't *that* your ultimate goal: to reduce your costs as much as possible long-term, while getting the job done?

IBM® System x® servers don't skimp on the essentials. They're designed from the ground up to be the most IT-friendly servers in the industry, incorporating those features that will give you the most productive, easiest to use, easiest to service, most available, and most cost-effective servers long-term. This paper describes how we take the best of industry-standard technologies—even helping to shape some of them through participation in standards groups—and combine them with IBM innovations to produce servers that offer you a unique set of capabilities with which to effectively manage your business.

IBM X-Architecture

The X-Architecture blueprint encompasses IBM System x and IBM BladeCenter® servers, IBM System x iDataPlex™ solutions, energy-efficiency and cooling via IBM Cool Blue™ technology, and proactive management.

It takes a comprehensive approach—using innovative hardware, software and services—to solve customer challenges today, and X-the Architecture blueprint embraces an evolving design methodology to address tomorrow's challenges. The fundamental *premise* of the X-Architecture blueprint will not change over time; however, the underlying *technologies* and how they are delivered will constantly evolve to meet the needs of our customers.

Gone are the days when IT personnel could throw a bunch of cheap servers in a room and call it a data center. Today, careful planning must be undertaken to validate the data center's ability to provide the necessary energy to servers, to cool them, and to grow over time without running out of room or overtaxing the data center's energy and thermal envelope.

Some long-term market trends continue, such as the ever-increasing need for performance. Meanwhile, new trends are emerging, including a greater concern over security and energy/thermal issues and a growing interest in virtualization. IBM has been visionary in recognizing these trends early and providing solutions via X-Architecture design to ready your business for the future.

Balanced System Performance

The x86 world has changed. Once, servers were divided clearly between scale-up and scale-out. While scale-out servers were relegated to running only single-threaded, single-processor applications—such as e-mail, and printing—the focus of monolithic scale-up systems was on running multithreaded, multiprocessor-capable applications, such as relational database managers.

Today, the line between scale-up and scale-out servers has blurred. Multi-core processors and large memory capacities allow uni- and dual-processor servers to run *scale-up* workloads that only a few years ago required 4-socket and larger servers. Performance becomes about utilization through virtualization, which is crucial to unlocking the value of multicore processors. It's an intelligent sharing of computing, storage, and information resources across different disciplines and departments in your organization. System resources are gathered into one pool—and can be allocated dynamically as needed—allowing servers and storage resources to be utilized more efficiently.

As the number of processor cores increase, system resources must keep up to avoid potential bottlenecks. For example, for 4-core processors to achieve maximum performance, they require four times as much memory as single-core processors, and as processors continue to increase throughput, so does memory.

IBM, utilizing its decades of experience in designing mainframes and multi-core supercomputers, leads in the design of x86 servers. IBM takes a holistic approach to balanced system design, optimizing its x86 servers to deliver outstanding performance and utility around the processors, memory, I/O, and network fabrics.

Innovation Comes Standard

Simply put, our philosophy regarding System x servers is “innovation comes standard.” We're delivering next-generation technology today that the competition can't match. Some server vendors view uni and 2-socket servers as commodities, using off-the-shelf components to produce cookie-cutter servers with no added value. Instead, IBM X-Architecture system design begins with standard parts and adds practical innovation to create something better: outstanding multi-core performance, high availability, scalability, energy efficiency, and proactive manageability. All aspects of the server's design are fair game for innovation.

Scalability

Normally, the term scalability doesn't appear in a discussion of uni- or dual-socket servers. That normally falls into the realm of 4-socket-and-up servers. But wouldn't it be nice to be able to buy a dual-socket blade server and then—if your needs change—magically turn it into a 4-socket server? Fortunately, no magic is required. There's the scalable IBM **BladeCenter LS42**. It starts out as a single-wide blade with up to two AMD Opteron processors, up to **32GB of 800MHz** DDR2 memory and one 2.5-inch HDD. Not enough? Simply snap an optional single-wide Multiprocessor Expansion Unit onto the first blade, in the adjacent blade slot. (No tools required.) This *doubles the capacity* of the now double-wide blade to **four** sockets, **64GB** of RAM and *two* HDDs. A snap-on **Storage and I/O Expansion Unit** blade supports up to three additional (hot-swap) HDDs. The LS42 also includes **four Gb Ethernet ports** and supports up to **12** communication ports total (via internal add-on cards). This sort of scalability for a dual-socket blade server is available from no one else.

Virtualization

Virtualization allows a uni- or dual-processor server (especially one with four cores) to act like several servers running separate operating system and application instances. Some customers focus on the total number of virtual machines they can run, when they should really focus on the amount of work being done in those virtual machines. If two servers support the same number of

virtual machines, the one with better performance will be doing more aggregate work. Thanks to our history, we understand how to build servers optimized for virtualization, including systems management tools to make the physical world aware of the virtual world. To truly excel in this virtualized world, your servers must have three things: 1) the capacity to support the needed number of virtual machines, 2) the throughput to perform well so the server can get more work done, and 3) the reliability to consolidate more work on that platform.

Normally, the process of setting up virtual machines on a server can take hours or even days. It involves installing the hypervisor software, downloading the latest updates, creating the virtual machines, and testing everything to make sure it's set up correctly. In order to simplify and speed up the process, specific models of the x3850 M2 ship with a **preinstalled, embedded hypervisor** (VMware ESXi) on an **internal flash memory drive** containing everything you need to have fully functioning virtual machines up and running in minutes. Similarly, all models of the x3550 M2 and x3650 M2 rack servers, and the HS22 blade server offer an optional flash drive with the same capability. All that's required is to plug the drive into the internal USB port.

Virtualization can expand the usefulness and extend the service life of your servers. But—and this is important—virtualization can require massive amounts of memory. If you have eight cores in your server but your application is single-core-only, you might consider running eight virtual machines (VMs) each with one instance of your application. But if your apps require 8GB of RAM per VM and your server is limited to 32GB, you're pretty much stuck with only four virtual machines per server and those other four cores go unused. By contrast, an IBM x3655 server would let you create eight VMs of 8GB apiece.

Memory

The trend going forward is to shoehorn ever more cores inside a processor. More cores can translate into faster overall processing performance, but only if other aspects of the server design don't act as bottlenecks to performance. For example, memory. Today's 4-core processors need four times the amount of memory that 1-core processors needed to generate the same sort of throughput per core. Having insufficient memory is a quick way to limit performance. Yet, many competitive systems do just that.

Many applications require massive amounts of memory for efficient processing. For those needs, we designed some of our rack servers (System x3550 M2, x3650 M2, and iDataplex dx360 M2) to support up to **128GB** of memory (16GB per core), a blade server (HS22) to support up to **96GB** (12GB per core), other servers (x3450 and x3655) to support up to **64GB** of memory (8GB per core) and the x3455 to support **48GB** (6GB per core). Not having enough memory capacity can strangle a memory-intensive application. Not only does having lots of memory benefit traditional "memory hogs," such as database servers, but with more cores comes the capability for virtualization.

Having more DIMM slots in your system offers you a second potential benefit: lower memory cost. Let's say you're running applications that don't require the maximum possible memory—you require only 32GB per system. In that case, it doesn't matter whether you buy an IBM server or one of the others, because they'll all support the same 32GB—right? Not really. With the competition offering only 8 DIMM slots, you are forced to buy all **4GB** DIMMs. If you've priced memory, you know that 4GB DIMMs cost more than double the price of 2GB DIMMs. With our x3450, x3550 M2, x3650 M2, and x3655 servers you can populate the 16 DIMM slots with lower-cost **2GB** DIMMs to achieve the same 32GB, saving you money. (You can also save with our x3455 and x3500 servers by using 8 2GB DIMMs and 4 4GB DIMMs.)

Memory capacity isn't the only way in which IBM servers offer you more performance than the competition. For example, in most AMD Opteron processor-based systems using 667MHz memory, if you populate all memory sockets, the speed of all DIMMs drops to only 533MHz.



However, the 1U System **x3455** offers **IBM Xcelerated Memory Technology™**. Because of this innovation, *all 48GB* in the x3455 run at the full 667MHz, unlike competitive systems.

Another superb performance option is the **x3450**, which matches **800MHz** fully buffered memory with a **1600MHz** front-side bus (FSB), up to a **3.4GHz 4-core Intel® Xeon®** processor, and an ultra-high-speed **x16 PCIe Gen 2** (16MBps) adapter slot.

Memory performance isn't everything, however. What good is top performance if your server crashes due to memory errors? To minimize downtime, we developed **IBM Chipkill™ ECC memory technology**, which is available in most System x servers¹. IBM-patented Chipkill protection offers up to 16 times the soft error correction of standard ECC memory.

Because protection against only soft memory errors isn't enough, we also offer two other levels of protection against hard errors: **memory mirroring** and **online hot-spare memory**. With mirroring, half the installed memory is reserved as backup to the other half². As an alternative, hot-spare memory takes a different approach. Memory is divided into either three or four sets of DIMMs (depending on the server). Each of the sets has one DIMM reserved as backup for the other DIMMs in the set. In this manner, only one-third or one-quarter of installed memory is reserved, leaving more memory available for processing. A number of System x servers offer mirroring³ and/or sparing⁴. Many competitive servers don't.

Storage

Another opportunity for innovation is in storage. Servers frequently require storage upgrades or servicing. The longer it takes to add or replace a drive, the higher the service cost. Most server vendors design their storage to be mounted in one of two ways: hot-swap or fixed. Each method has its pros and cons. **Hot-swap drives** offer the shortest servicing time, because you can add or replace a drive by pulling and inserting drive trays from the front of the unit without shutting down the system. However, hot-swap drives are expensive. As a result, servers offering hot-swap components cost more than those that offer only fixed drives. On the other hand, **fixed drives** are less expensive, but take longer to service because you have to shut down the server, open the system covers, undo a bunch of screws, and then reverse the process after installing the new drive. Plus, system downtime has costs of its own.

Wouldn't it be nice if there were a third alternative that offered the best of both worlds? In fact there is: **simple-swap drives**. An IBM-exclusive, simple-swap drives come on a drive tray much like those for hot-swap drives, allowing them to be inserted and removed in seconds. And they're much less expensive than hot-swap drives. They do require the server to be shut down, but because the servicing is faster, the actual downtime is far less than with fixed drives. In addition, many of our servers offer the flexibility of supporting *either* hot-swap or simple-swap drives⁵, in case your needs change. If hot-swap drives are more than you need, wouldn't you rather have quick-servicing simple-swap drives than the hassle of fixed drives?

Adding to the flexibility of drive *types*, IBM also offers the choice of drive *sizes* in many models⁶. The larger 3.5-inch drives offer higher capacity than 2.5-inch drives do, but the smaller drives allow the use of more physical devices in the server, including **12 drives** in the 2U x3650 M2 and **6 drives** in the 1U x3550 M2. This enables greater redundancy and



¹ Chipkill memory protection is offered in the x3400, x3450, x3455, x3500, x3550 M2, x3650 M2, and x3655 rack servers, as well as the HS12, HS22, LS22, and LS42 blade servers.

² In the x3550 M2 and x3560 M2, only two of the three memory channels are used for mirroring, with one channel mirroring the other.

³ Memory mirroring is available in the x3400, x3500, x3550 M2, x3650 M2, HS22, and the iDataPlex dx360 M2.

⁴ Online hot-spare memory is available in the x3400, x3455, x3500, x3655, and HS12.

⁵ Both drive types are supported on the x3200 M2, x3250 M2, x3350, x3400, x3550 M2, and x3650 M2 rack servers and the iDataPlex 2U chassis. The x3450 supports only simple-swap drives, while the x3500, x3610, and x3655 support only hot-swap drives.

⁶ The x3200 M2, x3250 M2, x3350, x3500, and x3655 rack servers and the iDataPlex 2U and 3U chassis all support *both* 3.5-inch and 2.5-inch drives.

speed via RAID arrays (RAID-60 in 2U servers; RAID-5/50/6 in 1U/2U servers). Also, 2.5-inch drives use approximately **40%** less⁷ energy than 3.5-inch HDDs, saving you a considerable amount of energy. (How many HDDs does your data center contain?) They also run cooler and require less airflow, permitting greater density.

Most servers offer a choice of high-performance Serial-Attached SCSI (**SAS**) hard disk drives or lower-cost, lower-performance Serial ATA (**SATA**) HDDs. IBM is no different, supporting either or both drive types in various servers. But we also support a third drive type in a number of IBM servers: 2.5-inch **solid-state drives (SSDs)**. These drives contain no spinning disks but consist of high-performance flash memory. Although of lower capacity than HDDs, these drives offer triple the reliability of spinning media⁸, and consume far less energy⁹. This makes them exceptional boot-to-SAN drives, as well as perfect for other uses that require small storage footprints but extremely high reliability. Additionally, they offer extremely fast read rates, so they're ideal for Web serving or other applications that are primarily read-only.

In addition, for those administrators concerned with the security of their data, most System x servers offer **removable tape media**¹⁰ internally. Other vendors require external media, potentially adding thousands of dollars in cost. As an alternative to internal tape storage, some System x servers also support the **IBM RDX USB Removable Hard Disk Drive**¹¹. Rather than using slow, cumbersome tape cartridges for backup, the DX drive uses removable high-speed 160/320/500GB disk cartridges, for the random-access advantages of disk storage and optional encryption. (External RDX drives are also available.)

Data security is a growing requirement for businesses of all sizes today due to security attacks, which are increasing in number as well as in sophistication. While many companies have invested heavily to thwart network-based attacks and other virtual threats, few effective safeguards have been readily available to protect against costly exposures of proprietary data at rest due to a system or hard drive being stolen, misplaced, retired, or redeployed. Compliance regulations require companies that have a data breach to inform customers of the breach, which can lead to a significant loss in their customer base. The IBM **ServeRAID™ -MR10is Vault** controller addresses these security concerns by combining advanced RAID functionality¹² with hardware-based IEEE-1619 XTS-AES 256 standard disk encryption. Until now, the only encryption option for x86 servers was software-based encryption. Software encryption requires servers to handle the encryption tasks in addition to data processing, which requires considerable processing power and results in overall degradation of server performance. But the ServeRAID controller offloads this burden from your system processors, providing data security without a performance penalty.

I/O

To achieve balanced system performance, you must match your I/O throughput to the performance of the rest of your system. There isn't much point in having ultra-fast processors and memory if the I/O creates a bottleneck.

To this end, IBM developed **eXtended I/O™** technology¹³ to provide the degree of flexibility and performance a dynamic infrastructure needs in I/O. eXtended I/O provides adapter slots on replaceable riser cards. This means you can mix-and-match PCI-X and PCIe slots as needed—

⁷ Seagate Technologies, "2.5-Inch Enterprise Disk Drives: the Key to Cutting Data Center Costs." <http://www.seagate.com/content/pdf/whitepaper/TP-534.pdf>.

⁸ MTBF of 73GB 2.5-inch 10K RPM HDD for the first 12 months = 305,167. MTBF of solid-state SanDisk SDD SATA 5000 2.5-inch 16GB = 918,298.

⁹ The 31.4GB solid-state drives (SSDs) offered in HS12, HS21, HS21 XM, HS22, LS22, and LS42 blade servers, as well as the x3550 M2 and x3650 M2 rack servers and the iDataPlex dx360 M2, consume only **2W** of power, compared with **9-10W** for 2.5-inch HDDs and up to **16W** for 3.5-inch HDDs. x3550 M2 and x3650 M2 servers also support 50GB SSDs.

¹⁰ Internal tape drives are supported on the x3100, x3200 M2, x3400, x3500, x3610, x3650 M2, and x3655.

¹¹ The internal RDX drive is supported in the x3200 M2, x3400, x3500, x3650 M2, and x3655.

¹² Supports RAID-0/1/10/5/50/6/60, with 256MB DDR2 battery-backed cache and ECC correction.

¹³ Servers featuring flexible eXtended I/O technology include the x3455, x3550 M2, x3650 M2, and x3655.

even HTx slots (in systems that use AMD Opteron processors). On various system x Servers we offer **x16 PCIe Gen 2** (16MBps) slots, **x16 PCIe** (8MBps) slots, or **HTx** (8MBps, with lower latency than x16 PCIe).

There are numerous ways to boost server performance via I/O beyond simply offering high-performance adapter slots. One is to offload work from the CPU to other processors. To this end, System x3450, x3550 M2, x3650 M2, and x3655 rack servers and HS22, LS22, and LS42 blade servers include Gigabit Ethernet controllers with **TOE** (TCP/IP Offload Engine). Similarly, most System x servers include an **integrated IBM ServeRAID** controller¹⁴. Not only do these controllers offload processing work from the CPU, but they contain large amounts of onboard cache memory to boost performance even more.

Another way to extend the I/O capability of 1U and 2U rack servers is to integrate controllers onto the motherboard, thereby freeing up valuable adapter slots. For example, most System x servers contain integrated systems management controllers (baseboard management controllers, or **BMCs**)¹⁵. Some also support an optional **IBM Remote Supervisor Adapter II Slimline** controller in a dedicated slot¹⁶ for additional systems management capabilities. Similarly, the x3550 M2 and x3650 M2 rack servers, the HS22 blade server, and the iDataPlex dx360 M2 all come standard with an **Integrated Management Module (IMM)**.

Reliability, Availability and Serviceability

Because all components have a finite life expectancy, it's crucial that your system not only offer many hot-swap and redundant components, but also that when service is needed, it can be done as quickly as possible to return your server to an operational state.

In addition to the aforementioned memory and disk features, System x servers offer many other features to warn of impending failure, diagnose problems, and minimize servicing time.

Predictive Failure Analysis (PFA) support provides advance warning that a component is *about to fail*. PFA affords you the opportunity to replace a component without data loss and potentially without downtime. Many competitive systems offer PFA only for memory, HDDs, and sometimes processors. On System x servers, PFA support is enabled for processors, memory, HDDs, power supplies, fans, and voltage regulator modules, depending on the server model.

In 1998, IBM introduced **light path diagnostics (LPD)**, a feature the competition has yet to match. Where most servers merely offer a system error LED on the front bezel, many System x servers¹⁷ offer a pop-out/drop-down diagnostic panel with an LED for each major component—processors, memory, etc. If a component fails, its diagnostic panel LED lights up. When the servicer opens the system cover, another LED beside the specific component—such as DIMM slot 12—identifies the failed part. Servicers don't have to waste time determining which of the 16 DIMMs needs replacement. And light path diagnostics works even when the server is unplugged and sitting on a workbench. (Pressing the Remind button will light the last LED illuminated.)



PFA works with LPD to turn on the appropriate light

¹⁴ Servers supporting integrated ServeRAID controllers include x3200 M2, x3250 M2, x3350, x3400, x3500, x3550 M2, x3650 M2, and x3655.

¹⁵ The x3350, x3400, x3450, x3455, x3500, x3610, x3655, HS12, LS22, and LS42 contain a full BMC. In addition, the x3200 M2 and x3250 M2 contain an integrated miniBMC.

¹⁶ The Remote Supervisor Adapter II SlimLine is compatible with x3200 M2, x3250 M2, x3350, x3400, x3500, and x3655.

¹⁷ Light path diagnostics is available in x3350, x3455, x3500, x3550 M2, x3650 M2, x3655, and all BladeCenter chassis.

path diagnostics indicator and optionally can send a message to the system administrator via the IBM Systems Director systems management software.

IBM Dynamic System Analysis (DSA) collects and analyzes system information to aid in diagnosing system problems. DSA also compares device driver and firmware on the system to the versions available on an UpdateXpress CD, providing a summary of the differences.

When servicing is inevitable, the ability to install and remove components quickly is crucial. To this end, we provide as many **tool-free components** as possible. All System x servers provide toolless covers, adapters, and DIMMs. Most servers offer hot-swap and/or simple-swap HDDs¹⁸. In addition, specific models also provide tool-free slide rails (for rack servers), optical drives, redundant hot-swap fans¹⁹ and power supplies²⁰, integrated RAID controllers, and integrated service processors.

Systems Management and System Support Features

IBM Systems Director is a platform management foundation that streamlines the way physical and virtual systems are managed across a multi-system environment. Leveraging industry standards, IBM Systems Director supports multiple operating systems and virtualization technologies across IBM and non-IBM platforms. Through a single user interface, it provides consistent views for visualizing managed systems and determining how these systems relate to one another while identifying their individual status, thus helping to correlate technical resources with business needs.

Base plug-ins provided with IBM Systems Director deliver core capabilities to manage the full lifecycle of IBM server, storage, network and virtualization systems. The base plug-ins include:

- **Automation Manager** — Performs actions based on system events
- **Configuration Manager** — Configures one or more systems resource settings
- **Discovery Manager** — Discovers virtual and physical systems and related resources
- **Remote Access Manager** — Provides a remote console, a command line and file transfer features to target systems
- **Status Manager** — Provides health status, alerts and monitors of system resources
- **Update Manager** — Notifies, downloads and installs updates for systems
- **Virtualization Manager** — Creates, edits, relocates and deletes virtual resources

Not only does IBM Systems Director work with system hardware, such as the BMC, Remote Supervisor Adapter II, and IMM, it uses an extensible modular interface that allows extensions and plug-ins to be added for additional functionality. Two of these are **IBM System Director Active Energy Manager™** and **IBM Systems Director Service and Support Manager** (previously called **IBM Electronic Service Agent™**). In addition, IBM Systems Director works with popular virtualization hypervisors, including **VMware**, **Microsoft Virtual Server**, and **Xen**.

IBM Systems Director Server is included with most System x servers²¹ for IBM AIX®, Microsoft® Windows®, Linux® on Power, Linux on x86, and Linux on System z.

Recognizing that often customers have large enterprise-wide frameworks in place, IBM Systems Director enables you to make the most of your existing enterprise management structure by upwardly integrating with these management software solutions:

¹⁸ Hot-swap and simple-swap drives are not available for x3100 and x3455.

¹⁹ Servers offering redundant hot-swap fans include x3400, x3500, x3550 M2, x3650 M2, x3655, and all BladeCenter and iDataPlex servers.

²⁰ Servers offering redundant hot-swap power supplies include x3200, x3250 M2, x3400, x3500, x3550 M2, x3650 M2, x3655, and all BladeCenter and iDataPlex servers. In addition, x3610 offers redundant (but not hot-swap) power supplies.

²¹ IBM Systems Director Server is provided with x3200, x3250 M2, x3350, x3400, x3455, x3500, x3550 M2, x3650 M2, x3655, and all BladeCenter and iDataPlex servers. The IBM Systems Director *Agent* is provided with x3100.

- Tivoli® Management Framework
- Tivoli NetView 7.1.x (Windows and Linux)
- Netcool/Precision IP (via SNMP)
- Netcool/Monitoring, Netcool/ISM
- Netcool/AEM (via SNMP)
- CA Unicenter NSM 3.1 and R11 (Windows)
- HP OpenView NNM 7.0.1 and 7.5.1 (Windows and Linux)
- HP OpenView Operations for windows 7.5x (Windows)
- Microsoft Systems Management Server 2003, Microsoft System Center Operations Manager 2007, and Microsoft Operations Manager 2005

Now you can manage your multi-system environment with the enterprise tools you already have in place.

IBM Systems Director Service and Support Manager 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 service²² if necessary to those customers entitled to onsite support under the terms of their warranty or an IBM Maintenance Agreement. Electronic Service Agent 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.

The Baseboard Management Controller (**BMC**), provided in many System x and BladeCenter servers, provides industry-standard **Intelligent Platform Management Interface (IPMI) 2.0**-compliant systems management²³. 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
- VLAN support
- Enhanced authentication and encryption algorithms (RMCP+, SHA-1, AES)
- Local update of BMC firmware
- Firmware firewall
- Support for IPMI v2.0 compliant management software (e.g., xCAT)
- Other mandatory and optional IPMI BMC functions

²² For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

²³ IPMI 2.0 support is provided with all System x servers except x3100, and all BladeCenter and iDataPlex servers.

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

In addition, the optional **Remote Supervisor Adapter II SlimLine** provides even more systems management capabilities, including:

- Graphical console redirection over LAN
- Web-based out-of-band control
- Windows “blue screen” capture
- Remote virtual floppy and CD-ROM
- High-speed remote redirection of PCI video, keyboard and mouse
- SSL (Secure Socket Layer) and LDAP (Lightweight Directory Access Protocol) support

The **Integrated Management Controller (IMM)** included in every x3550 M2 and x3650 M2 rack server, HS22 blade server, and iDataPlex dx360 M2 server, provides the function of both the BMC and the Remote Supervisor Adapter II SlimLine²⁴.

Unified Extensible Firmware Interface (UEFI), is the industry-standard replacement for the venerable and limited system BIOS. Used in PCs for several years, it is now moving into the server arena. UEFI is currently offered in the HS22 blade server, x3550 M2 and x3650 M2 rack servers, and the iDataPlex dx360 M2. This migration will continue into new servers over time. Some of the enhancements UEFI provides over the older BIOS, still used in most competitive servers, includes:

- 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
- Consistent firmware management across an entire product line

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*.

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 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.

The IBM **ServerProven**[®] program provides the confidence that specific options and operating systems have been tested on your System x server and are officially supported to work together. It is updated frequently and available online so that the latest compatibility information is always at your fingertips. Similarly, the IBM **ClusterProven**[®] program tests and identifies software as meeting the necessary high-availability requirements for the program.

²⁴ The x3550 M2 and x3650 M2 require an optional IBM Virtual Media Key to unlock the full advanced management capabilities of the IMM.

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 **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 Microsoft Windows Server 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. ServerGuide is included with all System x servers.

Energy Use and Thermal Efficiency

Managing energy and cooling resources has become a critical issue in the data center as energy consumption and thermal loads increase. Successfully meeting these energy and cooling challenges requires well-thought-out IT products, smart data center design, and a method for efficiently controlling and monitoring your systems' energy and heat requirements.

Taking this philosophy to heart, the IBM Cool Blue™ portfolio of tools demonstrates leadership thinking and sets the industry direction for optimal thermal efficiency.

IBM offers a number of technologies common to all System x servers to help improve your system and data center energy management:

- **IBM Systems Director Active Energy Manager** — In order to put control of processor energy-saving features at the fingertips of administrators, IBM developed Active Energy Manager, which is a powerful software tool, designed to take advantage of new processor features, such as balancing the performance of the system according to available energy input. A plug-in for IBM Systems Director, Active Energy Manager provides the ability to plan, predict, monitor and cap energy consumption based on your server configuration. Unlike competitive tools that merely add up the maximum potential power draw for each system—leading to grossly overblown estimates of total usage—Active Energy Manager uses embedded sensors to monitor *actual* energy use by component and server and uses that data to plot trends over time. In addition, it allows you to cap energy use by server or department, if needed. Active Energy Manager is available across the entire portfolio of IBM servers, from System x through System z.
- **IBM System x and BladeCenter Power Configurator** — This tool helps IT managers plan for data center power needs by providing the following information for specific configurations of System x and BladeCenter systems: energy input in watts, PDU sizing in amps, heat output in BTUs, airflow requirements through the chassis in CFM, VA rating, leakage current in mA, and peak inrush current in amps.
- **IBM Rear Door Heat eXchanger** — This is a sealed-coil water-cooled door that attaches to the back of an **IBM S2 42U Enterprise Rack**. It's designed to dissipate heat exiting the back of your servers before it leaves the rack. It can remove up to **50,000 BTUs** (14KV_a) per hour from the data center using chilled water lines under the raised floor. Improved cooling from the heat exchanger enables you to populate individual racks more densely using the same air conditioning, while freeing up valuable floor space.
- **Calibrated Vecteded Cooling** — Utilized in all System x, it provides extremely efficient cooling. This enables increased density of components without overheating. Innovations include using *tandem counterrotating fans*, angling fans to increase the cooling for specific components, and other features that promote greater airflow. (Four fans with counterrotating blades move more air and use less



energy than eight standard fans and mean fewer components that can fail.)

- Another thermal implementation, **isolated zone cooling**, requires only *some* fans in a specific “cooling zone” to switch to full speed in response to higher temperatures in that zone, rather than shifting *all* of them into high gear. Because of this design, the fans use less energy and create less ambient noise than other servers.
- The servers all use **hexagonal ventilation holes** in the chassis. Hexagonal holes can be grouped more densely than round holes, providing greater airflow through the system cover.
- The **x3550 M2** and **x3650 M2** also include an **altimeter**. At high altitudes, the air is thinner and doesn’t cool as well as at lower elevations. In most servers, the fans run fast all the time to allow for use at high elevations, wasting power at lower elevations. The altimeter allows the IBM servers to run at reduced speed at lower altitudes.

These innovations help you to save you money on energy costs, put less wear-and-tear on the fans, and produce less ambient noise in the data center.

- **Energy-smart power supplies** — The typical power supplies used in the server industry are approximately **70-75%** efficient at converting AC wall current into the DC power used inside a server. By contrast, the power supplies IBM uses in System x servers and BladeCenter chassis are significantly more efficient—up to **92%** efficient in some cases. This helps save you money both on energy consumption up front and on cooling at the back end.
- **Low-voltage processors** — Intel offers low-voltage versions of some processors, which run at the same clock rates as their higher-voltage cousins but consume less energy. Not all server vendors offer these low-voltage processors, which costs you money in wasted energy and cooling. For example, we offer a **60W** (2.26GHz or 2.13GHz) 4-core Intel Xeon processor in the x3550 M2 and x3650 M2 rack servers, as well as the HS22 blade server and the iDataPlex dx360 M2. For less-demanding workloads, we also offer a **35W** (2GHz) 1-core Intel Celeron® D processor in the x3250 M2. Many of our servers also offer **65W** 2-core Xeon or Core™ 2 Duo processors.
- **Energy-smart memory** — Many System x servers use energy-efficient unbuffered **DDR2** memory²⁵, which consumes up to **37% less energy** than competitive systems using Fully Buffered DDR2 DIMMs. The x3550 M2, x3650 M2, the HS22, and the dx360 M2 all use registered **DDR3** memory which is **10-15%** more energy efficient than even DDR2 memory.
- **2.5-inch HDDs** — 2.5-inch drives use approximately **40% less**²⁶ energy than 3.5-inch HDDs, saving you a considerable amount of energy. (How many HDDs does your data center contain?) They also run cooler and require less airflow, permitting greater density. 2.5-inch drives are available in x3200 M2, x3250 M2, x3350, x3500, x3550 M2, x3650 M2 and x3655 servers, as well as all BladeCenter and iDataPlex servers.
- **Solid-state storage** — The 31.4GB solid-state drives (SSDs) offered in HS12, HS22, LS22, and LS42 blade servers, as well as the x3550 M2 and x3650 M2 rack servers and the iDataPlex dx360 M2, consume only **2W** of power, compared with 9-10W for 2.5-inch HDDs and up to 16W for 3.5-inch HDDs. Not only does this save even *more* power, but SSDs, with no moving parts, offer triple the reliability of spinning media²⁷.

The innovative and efficient cooling designs of System x and BladeCenter have helped make us more energy-smart than the competition. For example, **BladeCenter E** uses up to **58%** less power than Dell 1U servers (with a similar reduction in excess heat production) and up to **32%** less power than HP 1U offerings²⁸. IBM **BladeCenter H** requires nearly **10%** less power than the

²⁵ Unbuffered DDR2 memory is used in x3100, x3200 M2, x3250 M2, x3350, x3455, x3610, and x3655 servers.

²⁶ Seagate Technologies, “2.5-Inch Enterprise Disk Drives: the Key to Cutting Data Center Costs.” <http://www.seagate.com/content/pdf/whitepaper/TP-534.pdf>.

²⁷ MTBF of 73GB 2.5-inch 10K RPM HDD for the first 12 months = 305,167. MTBF of solid-state SanDisk SDD SATA 5000 2.5-inch 16GB = 918,298.

²⁸ IBM BladeCenter H with HS21, 10 servers, 20 processors = 2,703W and 9,220 BTUs. Dell 1950 1U, 10 servers, 20 processors = 4,276W and 14,584 BTUs. HP DL 360 G5 1U, 10 servers, 20 processors = 3,570W and 12,173 BTUs. IBM power engineering test data. Numbers are average worst-case for P6 Burn exerciser program, tested in the IBM lab. Blade power is the average power of total chassis solution.

equivalently configured HP BladeSystem c-Class 7000²⁹. Extrapolated over 224 servers and with an energy cost of 9.4 cents per kilowatt hour, this can save you up to \$12,000 per year per chassis.

Summary

In today's economy, Everyone is looking to slash costs. IBM System x server innovations can save you money in many ways:

- By reducing servicing time via tool-free components, Dynamic Systems Analysis, Predictive Failure Analysis, and light path diagnostics
- By allowing the use of lower-capacity/lower-cost memory to achieve the same total memory capacity as servers with fewer DIMM slots
- By keeping your servers online longer through the use of redundant, hot-swap, and simple-swap components and solid-state drives
- By offering simple-swap drives as a lower-cost alternative to hot-swap drives, and a faster service alternative to fixed drives
- By using energy-smart components, such as low-voltage processors, 92% efficient power supplies, counterrotating fans, altimeters, 2.5-inch HDDs and SSDs, and DDR2/DDR3 memory
- By reducing virtual machine setup time via an embedded hypervisor
- By giving your administrators remote access to servers—thus cutting travel time
- By enabling greater longevity for your servers, through greater memory capacity and virtualization
- By providing no-cost/low-cost tools to let you configure, monitor, and even cap energy use
- By offering products to keep your data center cooler, thus helping to reduce your cooling costs

To put it all in perspective, consider the following scenario: You buy an IBM rack server. The toolless slide rails and cable management arm make installing the server and neatly organizing the cables a piece of cake. Later, to add a new HDD, simply pull out the hot-swap drive cover, slip in the HDD and close the locking handle. All this takes only seconds. Let's say that somewhere down the road a fan's RPMs drop below the acceptable threshold. Immediately PFA catches this and, working with the BMC/IMM and IBM Systems Director, alerts the server administrator by email or text message to the potential impending failure. The administrator visits the server, checks the pop-out light path diagnostics panel, opens the cover, looks for the LED that identifies the specific hot-swap fan in need of service, and replaces it while the system is still up and running. Again, all this can be done in seconds and without tools. Also, because our servers use less energy and produce less heat than the competition, and because we offer two Rear Door Heat eXchanger models (one for standard racks and one for iDataPlex racks)—which can remove great quantities of heat from your data center—it can cost far less to run and cool your data center. With a lower ambient temperature in the data center, your servers might even last longer.

To paraphrase Benjamin Franklin, when it comes to servers time saved is money earned. System x servers can save you a lot of time and a lot of money through innovation.

²⁹ Edison Group Blade Server Power Study, *IBM BladeCenter and HP BladeSystem*, November, 2007.



For More Information

IBM System x Servers

IBM System x and BladeCenter Power Configurator

IBM Systems Director Active Energy Manager

IBM Standalone Solutions Configuration Tool (SSCT)

IBM Electronic Service Agent

IBM ServerProven Program

IBM Technical Support

IBM Configuration and Options Guide

ibm.com/systems/x

ibm.com/systems/bladecenter/powerconfig

tinyurl.com/3lafyd

ibm.com/servers/eserver/xseries/library/configtools.html

ibm.com/support/electronic

ibm.com/servers/eserver/serverproven/compat/us

ibm.com/server/support

ibm.com/servers/eserver/xseries/cog

Legal Information

© IBM Corporation 2008

IBM Systems and Technology Group

Dept. U2SA

3039 Cornwallis Road

Research Triangle Park, NC 27709

Produced in the USA

March 2009

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, the e-business logo, Active Energy Manager, AIX, BladeCenter, Calibrated Vectored Cooling, Chipkill, Cool Blue, ClusterProven, Electronic Service Agent, eXtended I/O, iDataPlex, Predictive Failure Analysis, ServeRAID, ServerGuide, ServerProven, System x, and Xcelerated Memory Technology, 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 on the Web at <http://ibm.com/legal/copytrade.shtml>.

Intel, Intel logo, Celeron, Core, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft 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.

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.