



IBM X-Architecture Blueprint 2009: An Overview

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

Executive Overview

A decade ago, our biggest IT concerns were such things as processor frequency, memory and storage capacity and network bandwidth. Today we're more concerned with improving productivity through virtualization, optimization, energy stewardship, and flexible sourcing. We need to provide high availability and quality of existing services, but also meet customer expectations for real-time, dynamic access to innovative *new* services. And somehow we must find a way to not only address today's security, resiliency, and compliance challenges, but also prepare for the new risks posed by an even more *connected* and *collaborative* world. But how do we accomplish all this?

Fortunately, IBM foresaw the need for flexible, expandable servers that could evolve with market demands to help reduce costs, manage risk, and improve service—what we now call dynamic infrastructure. In 1998, recognizing that x86 servers were no longer mere commodities and needed many of the same reliability, serviceability, availability, and performance attributes as mainframes and other advanced servers, IBM® introduced x86 servers designed with the IBM X-Architecture® blueprint. This game-changing design philosophy incorporated x86 industry firsts such as Chipkill™ memory protection, light path diagnostics and Predictive Failure Analysis®, and helped elevate x86 servers from their previously limited roles to their present status as mission-critical, mainstream enterprise servers. The competition has been playing catch-up ever since.

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.

The X-Architecture blueprint takes a comprehensive approach to providing a dynamic infrastructure—using innovative hardware, software and services—to solve customer challenges today, and it embraces an evolving design approach to address tomorrow's challenges. 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 provide the greatest possible flexibility in System x, BladeCenter, and iDataPlex servers. 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.

This paper describes how the melding of industry-standard components with IBM innovation produces servers offer you a unique set of capabilities with which to effectively manage your business. You'll see why IBM systems designed with the X-Architecture blueprint run faster—yet cooler—use less energy, and are more reliable, easier to manage and better for virtualization than competitive servers, *and* offer a lower total cost of ownership (TCO).

Market Trends

Gone are the days when IT personnel could throw a bunch of cheap servers in a room and call it a data center. In today's data center, careful planning must be undertaken to validate its 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. It also requires a dynamic infrastructure—one that's able to react intelligently and effectively to change, is secure, is cost-efficient, and is Green.

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. More and more companies are looking to "go green," not only because it's good for the planet, but also because it's good for the

bottom line. 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 into scale-up vs. 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 allow *all* servers to run *scale-up* workloads to some extent. At the same time, scale-up servers can run multiple virtualized *scale-out* workloads. Sure, we all want to reduce costs and Go Green, but we still need top performance, because getting the job done quicker is one way to cut expenses. Therefore, performance is less about processor frequency than it is about better utilization of all aspects of the system. We want to use every bit of that expensive processor and all of the memory we paid for. Performance is now all about utilization through virtualization, which is crucial to unlocking the value of multicore processors. It's an intelligent sharing of computing, storage or 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 six-core processors to achieve maximum performance, they require six times as much memory as single-core processors. IBM, utilizing its decades of experience in designing virtualized 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.

Enterprise Rack Servers

Using the mainframe-inspired IBM-designed eX4 chipset, the **x3850 M2** and **x3950 M2** with XpandOnDemand™ modular expansion capability are the only switchless x86 servers from a Tier-1 vendor capable of scaling to **16** sockets and **96** processor cores, **1TB** of RAM, **1GB** of L4 cache, and **28** PCIe slots (including **8** hot-plug PCIe slots) in only 16U of rack space.



The **x3850 M2** is a 4U 4-socket scalable rack server that supports 6-core Intel® Xeon® processors, up to 256GB of DRAM, 256MB of L4 cache, seven high-speed adapter slots (including two hot-plug slots) and four HDDs.

If you ever find the need for more resources, simply install the ScaleXpander Option Kit to turn it into an **x3950 M2**, then attach another x3950 M2 chassis to the first using an XpandOnDemand cable. It immediately doubles your processors, memory, adapter slots, HDD bays and L4 cache capacity. Still not enough? Add up to two more x3950 M2 chassis. No other x86 server on the market offers this level of flexibility, scalability, or performance.

Due to the design of AMD Opteron 8000 Series processors, 3-processor configurations produce suboptimal performance out of the box. Because of this, the competitive servers force you into 2- and 4-processor configurations—whether that's what you want or not. In another example of System x design flexibility, IBM took a different approach for the **x3755**: Instead of limiting your options, we chose to *solve the problem*. The result is the IBM **CPU Pass Thru Card**. Installed in the unused or fourth CPU/Memory card slot, it enables full performance in all configurations. A **3-processor** configuration offers you the option of *higher performance* than a 2P system, and *lower cost* than a 4P configuration—with the ability to expand later. In fact, for some applications a 3P

solution (including the CPU Pass Through Card and built-in IBM **Xcelerated Memory Technology**[™]) will outperform the competition's 4P servers.

The Opteron processor supports DRAM speeds of up to 667MHz. However, due to the way the processor is designed, in most systems when the DIMM slots are fully populated, the speed of *all* system memory drops to only 533MHz. The System x3755 avoids this limitation through Xcelerated Memory Technology. This innovation allows all RAM in the x3755 to run at the full 667MHz, using inexpensive industry-standard DIMMs. This gives x3755 memory as much as a **15%** performance boost over the competition.

Rack Servers

Simply put, our philosophy regarding rack-based 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 quad-core and six-core performance, high availability, scalability, energy efficiency, and proactive manageability. For example:


- The IBM **System x3650 M2**, a 2U server, supports up to **12 2.5-inch hot-swap** high-performance **SAS** HDDs or high-reliability/low-energy-draw **solid-state drives**. For those concerned with the security of their data—and who isn't?—the **x3650 M2** is the only 2U server from a Tier-1 vendor to offer **removable tape media** internally. Other vendors require external media, adding thousands of dollars in cost. In addition, the **x3650 M2**, along with the **x3550 M2** and other System x servers, support the IBM **ServeRAID-MR10iis Vault**, the industry's only hardware-based data encryption offering available for x86 servers. (For more information on this controller, see the **Virtualization and Security** topic, below.)
- Because of Xcelerated Memory Technology, *all* 48GB in the **System x3455** runs at the full 667MHz, unlike competitive Opteron processor-based systems that have to drop to 533MHz when fully populated with memory.
- Another superb performance option is the **System x3450**, which matches **800MHz** fully buffered memory with a **1600MHz** front-side bus (FSB), up to a **3.4GHz quad-core** Xeon processor, and an ultra-high-speed **x16 PCIe Gen 2** (16MBps) adapter slot.
- An innovative feature used in several System x servers is **IBM eXtended I/O**[™] technology. IBM understands that 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 ultrafast processors and memory if the I/O lags behind. Although tower servers have plenty of room to accommodate different types of adapter slots, 1U and 2U rack servers are another story. With eXtended I/O, you needn't worry about getting stuck with the wrong configuration. eXtended I/O provides adapter slots on replaceable riser cards. This means you can mix-and-match PCI-X and PCIe slots as needed—even HTX slots (in our Opteron-based 1U and 2U servers).



BladeCenter Chassis

Unlike the competition, IBM doesn't have a one-size-fits-all mentality. For the ultimate in flexibility, IBM offers *five* different blade chassis:

BladeCenter S (with integrated mass storage, 110V power, and office-friendly acoustics) for small and medium businesses or departmental needs. **BladeCenter E** for typical data center requirements, and **BladeCenter H** for higher-performance needs, supporting 10Gb fabrics. For telecom environment, we offer **BladeCenter T**, a ruggedized NEBS 3/ETSI-certified platform, and **BladeCenter HT**, a higher-performance version supporting 10Gb fabrics.



For a non-data-center office environment, BladeCenter S has integrated disk storage with up to **12TB** of SATA or **5.4TB** of SAS storage and—with the optional SAS RAID Controller—enables a SAN in the chassis. HP has to give up blade space to accommodate more disk storage. The **IBM BladeCenter S Office Enablement Kit** is a wheeled enclosure that easily slides under a standard office desk. It includes a locking door and an acoustic dampener to keep noise to an acceptable level. HP does not offer this. An optional attachment filters out dust and dirt.

In addition to the high-speed switches—including 10Gb Ethernet, 4X InfiniBand[®], and 8Gb Fibre Channel—supported by BladeCenter H and BladeCenter HT, all BladeCenter chassis support a number of 1Gb switch or bridge modules and fabrics, including Gigabit Ethernet, iSCSI, and SAS. Due to the high degree of integration in the chassis, all blades have two 1Gb paths to the switches and bridges natively. By adding a daughter card to each blade, two or four *additional* 1Gb or two 10Gb ports can be added. Using an optional **IBM Multi-Switch Interconnect Module** (MSIM) in a BladeCenter H or BladeCenter HT high-speed switch bay *doubles* the number of ports available to *any* blade server in the chassis.

QLogic Ethernet and **Fibre Channel InfiniBand Bridge** modules in conjunction with the **Cisco 4X InfiniBand Switch** provide gateway functions for one or more InfiniBand-connected groups of BladeCenter H or BladeCenter HT systems to external Gb Ethernet or Fibre Channel networks. This further **lowers the cost of ownership** from an acquisition cost and total cost of ownership perspective.

To further advance next-generation thinking around blades, IBM formed **Blade.org**. It supports openness and collaborative innovation around new products and systems. Blade.org has invested more than \$1 billion in the last few years to fund companies developing emerging technologies and solutions to simplify computing with blade servers. Membership includes leaders in blade software and hardware, developers, distribution partners and end users. Collectively, the community has introduced dozens of solutions to the market to help customers choose with confidence from the blade portfolio. We currently offer networking switches from QLogic, Brocade, Cisco and BNT. Whether you're looking for low-cost Level 2 Ethernet, advanced Level 2/3 or Level 2-7 Ethernet, 10Gb Ethernet, 4Gb Fibre Channel, 4X InfiniBand, SAS or iSCSI, BladeCenter offers them all. You won't be locked out of the latest-and-greatest technology with BladeCenter.

BladeCenter Blade Servers

When it comes to blade servers, IBM has no equal. The competition offers dual-socket server blades containing either Intel or AMD processors. We offer single-, dual-, and quad-socket servers, a choice of Intel, AMD, IBM POWER6[®], or Cell Broadband Engine[™] processors, even a workstation blade—and they can all be mixed-and-matched within a single chassis.

The new **IBM BladeCenter HS22** supports up to **two 2- or 4-core Xeon 5500-series** processors, which offer up to **225%** better performance than the previous generation (and up to **900%** better performance than the single-core processors of three years ago that you may still be using), up to **96GB** of memory, running at 800-1333MHz, a choice of two 2.5-inch **hot-swap SAS or SATA** HDDs or **hot-swap solid-state** drives, and planned **NEBS-3/ETSI-compliance**¹. It also supports optional embedded virtualization on an internal USB flash drive.

On a budget, or have workloads that can't take advantage of multiple processors anyway? The **IBM BladeCenter HS12** offers a single processor, up to **24GB** of memory, a choice of two 2.5-inch HDDs or **solid-state** drives, and does so at a lower cost than competitive two-socket blades.

If your needs run to *scalable* blade servers, there's the scalable **IBM BladeCenter LS42**. Two AMD Opteron processors, up to **32GB** of **800MHz** DDR2 memory and one 2.5-inch HDD aren't

¹ Selected models.


enough? Simply snap an optional Multiprocessor Expansion Unit onto the first blade in the adjacent blade slot. 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).

Are you an IBM System i™ or System p™ user? The single-socket **IBM BladeCenter JS12** and dual-socket **JS22** blades can run the IBM i™, IBM AIX®, and Linux® operating systems and applications. How about high-performance computing? With two **IBM PowerXCell™** (enhanced Cell/B.E.) processors and up to 32GB of memory, the **IBM BladeCenter QS22** blade offers the ultimate in HPC.


iDataPlex Solution

IBM designed the System x iDataPlex solution to address the needs power, cooling, and high density needs of large data centers. It takes standard rack server design and stands it on its ear. Combining the best features of rack and blade servers, and adding a few new twists, we managed to design racks with **more than double** the compute density of standard racks (up to 102U of servers, storage, and switches) and up to **five times** the compute density of the data center, while using less power and producing less waste heat.

The iDataPlex solution incorporates the best of both rack and blade server design, by combining 2U and 3U chassis with interchangeable server, I/O, and storage trays that can be inserted and removed quickly and easily. The same **2U Flex Chassis** can contain either **two 1U server nodes, one server node and one I/O tray, or one server node and one storage tray**. The **3U Chassis** supports **one server node** along with **2U of storage**, for as many as a **dozen** drives (12GB SATA or 5.4GB SAS) and a server in only 3U. Storage options for these chassis include combinations of **3.5-inch SAS or SATA drives, 2.5-inch SAS, SATA, or solid-state drives**, and a choice of **hot-swap or simple-swap**. There are *nearly a dozen possible configurations* of chassis and trays supported.



This does not even count the choice of three different server trays that can be used: the **dx320**, the **dx340**, and the **dx360 M2**. They offer a wide range of processor options (including **two-core** and **four-core; standard- and low-voltage**, Xeon **5200, 5400, and 5500** series), as well as memory options (registered DDR2, fully buffered DDR2, and DDR3). This versatility offers you outstanding configuration flexibility, pricing flexibility, and investment protection. For example, if your needs change you can easily swap out trays to convert an I/O Server configuration into a Compute Server configuration or a Storage Server configuration, or vice versa, as well as upgrade from a dx320 server node to a dx360 M2 node. A few of the possible configurations include:



In addition, by using the optional **IBM Rear Door Heat eXchanger for iDataPlex**, you can remove *more than 100%* of the heat emitted by the servers in the rack, actually helping to *cool* the data center. Because of this, the iDataPlex solution permits data center densities up to 5X that of data centers using conventional servers and racks. And to simplify installation, everything is ordered preconfigured and preinstalled in racks, shipped that way, and then set up for you in your data center. (For more information about the iDataPlex solution, contact an IBM sales rep or authorized business partner.)

Availability Features

High performance is important, but only when a server is up and running. With virtualization enabling consolidation of many servers onto one, reliability, availability and serviceability are more important than ever. To this end, IBM incorporates a number of innovations as standard features in most System x and all BladeCenter servers.

IBM Memory ProteXion™ technology works in conjunction with **IBM Chipkill** technology and standard ECC protection, to provide *three-level* memory correction.

In 1998, IBM introduced **light path diagnostics**, a feature the competition has yet to match. Where most servers merely offer a system error LED on the front bezel, IBM offers 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 32 DIMMs needs replacement. And light path diagnostics works even when the server is unplugged.



Extensive **Predictive Failure Analysis (PFA)** support provides advance warning that a component is *about to* fail. It illuminates the appropriate light path diagnostics indicator and optionally can send a message to the system administrator via the IBM Systems Director systems management software². On System x enterprise servers, PFA support is enabled for processors, memory, hard disk drives, power supplies, fans, voltage regulator modules, and the Xcel4v Dynamic Server Cache. By comparison, many competitive systems offer PFA only for memory, HDDs, and sometimes processors.

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.

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 both System x and BladeCenter to help improve your server and data center energy management:

- **IBM Rear Door Heat eXchanger** — This is a 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 (14KVa) per hour from the

² Provided as a standard feature on most servers.

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. (The **Rear Door Heat eXchanger for iDataPlex** is a double-wide version of this, removing up to 100,000 BTU.)

- **Calibrated Vectored Cooling™** — Utilized in System x and BladeCenter servers, it provides extremely efficient cooling. This enables increased density of components without overheating. Innovations include using *tandem counterrotating fans* (which can move more air with less energy than two standard fans can), angling fans to increase the cooling for specific components, *hexagonal ventilation holes* (which can be grouped closer together than round holes), and other features that promote greater airflow. 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.
- 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. The altimeter allows the IBM servers to run at lower speeds 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 the case of **x3550 M2**, **x3650 M2**, **BladeCenter H**, and **BladeCenter E** with the new higher-efficiency power supplies. This helps save you money both on energy consumption up front and on cooling at the back end.
- **Low-voltage processors** — Intel, AMD, and IBM offer 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.
- **Low-power memory** — eX4 enterprise servers use Buffer on Board technology with DDR2 memory, which can save over 200 watts per hour over competitive systems that use Fully Buffered DIMMs. In addition, the **x3550 M2**, **x3650 M2**, **HS22**, and iDataPlex **dx360 M2** servers all use **DDR3** memory, which is **10-15%** more power-efficient than even DDR2 memory.
- **2.5-inch HDDs** — 2.5-inch drives use approximately **40%** less³ 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.
- **Solid-state storage** — 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. Not only does this save even more power, but SSDs, with no moving parts, offer triple the reliability of spinning media⁴.

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 offerings⁵. IBM BladeCenter H requires nearly 10 percent less power than the

³ 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>.

⁴ MTBF of 73GB SFF 10K RPM for first 12 months = 305,167. MTBF of solid-state SanDisk SDD SATA 5000 2.5 16GB = 918,298.

⁵ 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 7000⁶. 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.

Similarly, in vConsolidate⁷ benchmarks the 4P x3850 M2 and 8P x3950 M2 consumed **13.6% less** energy than the HP DL580 G5, and **19.6% less** energy than two 4P HP DL580 G5 servers, respectively. (HP doesn't offer an 8-processor x86 server.) Comparing performance-per-watt, the x3850 M2 had **29% better** performance-per-watt than the HP DL580 G5, and the x3950 M2 had almost **36% better** performance-per-watt than the 4P HP DL580 G5.

IBM Systems Director Active Energy Manager

In order to put control of processor energy-saving features at the fingertips of administrators, IBM developed **IBM Systems Director Active Energy Manager™**. Active Energy Manager (AEM) 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 (described below) AEM provides the ability to plan, predict, monitor and cap energy consumption based on your System x or BladeCenter hardware configuration.

IBM Power Configurator

The **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: 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.

Virtualization and Security

Even the smallest servers become SMP-capable via multi-core processors. Because of this, IBM is investing in system and data center management technologies to help you achieve higher utilization levels. Having better control allows you to proactively reassign workloads to other assets and temporarily shift bandwidth as needed. This can help you delay the need for more servers, manage electricity usage, assign failover servers for improved redundancy, and so on.

Virtualization means more applications per server, which puts more demands on the system. As a result of virtualization, the focus of high-end scalable x86 systems has shifted from a hardware-centric view to one of flexibility and cost-per-virtual-machine. Running 20 or 50 or 100 applications on the same server instead of one task per server is far more cost-effective. And it allows the applications to share a dynamically allocated pool of common system resources for improved efficiency.

IBM offers servers that are extraordinarily well suited for virtualization. For example, the x3850 M2 and the BladeCenter HS21 XM now offer selected models *integrated* with **VMware ESXi** preloaded on an **internal flash drive**. Similarly, the HS22 blade server offers an optional upgrade (for all models) to VMware ESXi preloaded on an internal flash drive. The integrated hypervisor 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.



Virtualization, naturally, doesn't involve only hardware. One tool, **IBM Virtualization Manager** allows you to manage physical and virtual machines from a single console. Another offering, **IBM BladeCenter H and Cisco VFrame Solution**, provides fabric sharing and I/O consolidation, scales to **512** servers—far more than competitive offerings—and is simple to use. It virtualizes the server, workload and network connectivity, offering policy-driven automated blade failover. **IBM BladeCenter Address Manager** allows for the I/O virtualization of Ethernet and Fibre Channel connections *within a system* by providing users the option to assign Ethernet and Fibre Channel

⁶ Edison Group Blade Server Power Study, *IBM BladeCenter and HP BladeSystem*, November, 2007.

⁷ <http://www.principledtechnologies.com/clients/reports/IBM/IBMvCon0808.pdf>. As of August 27, 2008.

port addresses used by their server blades via software as an alternative to the addresses that are burned in to the hardware during manufacturing. IBM **BladeCenter Open Fabric Manager** allows you to deploy servers and I/O in minutes, rather than hours and it provides automated blade failover (inter- and intra-chassis) for up to 100 chassis. It's supported across all BladeCenter chassis, blades, and switches. If you use third-party management software for your networking equipment, you can seamlessly leverage the I/O virtualization capabilities that Open Fabric Manager delivers without disrupting your current management structure.

The increased spread of SANs, both Fibre Channel and iSCSI, is driving the need for **storage virtualization** as well. IBM storage virtualization tools can help transform the economics of enterprise storage by enabling users to simplify their infrastructure, protect their data and efficiently manage information. IBM **System Storage™** tools allow you to virtualize disk and tape storage. Solutions include the **System Storage SAN Volume Controller**, **SAN File System**, **IBM Tivoli® Intelligent Orchestrator** and **IBM Tivoli Provisioning Manager**.

Security concerns will eclipse TCO as the driver of *recentralization* and server consolidation will spread to desktops via virtual clients. To this end, IBM's **Virtual Client solution** allows users to enjoy all of the benefits and personal control of a stand-alone desktop—including print capabilities, USB drive support, and audio—while reducing many of the challenges related to current stand-alone desktop environments. These include limiting susceptibility to theft and viruses, ease of deployment of new users, extended downtime during a hard drive failure, or having to rebuild their preferences and settings after each client “refresh.”

Data security is a growing requirement for businesses of all sizes today. 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 potentially costly exposures of proprietary data due to a system or hard drive being stolen, misplaced, retired, or redeployed. The IBM **ServeRAID-MR10is Vault** controller addresses these security concerns by combining advanced RAID functionality⁸ 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.

Proactive Management

Inexpensive servers aren't inexpensive if they require excessive administrative or servicing time to bring them online and keep them that way. Labor is becoming the most expensive element in server deployment and ownership. Tools that simplify deployment, monitor system health, limit energy usage and report pending problems, are crucial to keeping costs under control.

IBM offers a number of tools to help you tame the complexity of systems management and administration, while managing costs:

- **IBM Systems Director** for advanced workgroup management is included with all BladeCenter and most System x servers. From a single user interface, IBM Systems Director enables monitoring and event management of physical and virtual resources across a heterogeneous IT environment. If problems arise, IBM Systems Director can send alerts to administrators via e-mail, cell phone, pager and other means. And unlike the competitors' offerings, IBM Systems Director is vendor neutral and supports a number of compatible non-IBM systems.
- **IBM Continuous Availability Manager (CAM)**, an extension for IBM Systems Director, manages the availability of virtual containers, virtual machines, hypervisors and physical hardware. It can predict unplanned outages and react to prevent or reduce their duration. Its “hypervisor-neutral” design allows it to work equally with Xen, VMware, Microsoft Virtual

⁸ Supports RAID-0/1/10/5/50/6/60, with 256MB DDR2 battery-backed cache and ECC correction.

- Server, or none at all. CAM requires no internal agents and supports multiple open interfaces.
- **IBM Virtualization Manager** (described previously) is another optional extension for IBM Systems Director.
 - **IBM Systems Director Active Energy Manager** (described previously) is a no-cost extension for IBM Systems Director that allows the administrator to plan, predict, monitor, and even limit energy consumption by System x and BladeCenter servers.
 - **IBM Remote Deployment Manager (RDM)**, an IBM Systems Director module, supports the deployment, update and retirement of servers, personal computers, workstations and point-of-sale terminals across the computing environment.
 - **IBM UpdateXpress™**, a no-charge Web-based tool, can help reduce your cost of computing by providing a simple yet effective way for administrators to update the server or client system firmware from anywhere on the network.
 - **IBM Integrated Management Module (IMM)** combines both the diagnostic functions of the integrated baseboard management module (**BMC**) used in earlier IBM servers and the advanced systems management functions of the optional **Remote Supervisor Adapter II** available for most System x servers. The IMM comes *standard* in the **HS22** blade server, **x3550 M2** and **x3650 M2** rack servers, and the iDataPlex **dx360 M2** today, and will be included in future servers as well.
 - **IBM ToolsCenter** consolidates the needed tools for managing servers individually into an integrated suite. The tools are organized by function: deployment, updates, configuration, and diagnostics. Tools are now simpler to access and use with a single easy-to-use 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.
 - Integrated **Trusted Platform Module (TPM) 1.2** is support in selected servers.
 - **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

Support and Services

IBM has been renowned for decades for its award-winning, world-class technical support and services. For example, in October 2007 IBM received the prestigious Service & Support Professionals Association (SSPA) STAR award for Service Excellence in Continual Improvement for System x, BladeCenter and IntelliStation. The award is presented to the company that demonstrates exceptional service levels and customer satisfaction for three or more years, with year-over-year improvements and a plan in place for continued future improvement.

Some of the innovative services we offer, to help you optimize and maintain your BladeCenter and System x solutions and overall data center, include:

- **Asset Recovery Solutions**
- **Data Center Consolidation and Relocation**
- **Data Center Energy Efficiency Assessment**

- **Data Center Stored Cooling Solution (“the cold battery”)**
- **High Density Computing Data Center Readiness Assessment**
- **High Density Computing Integrated Rack Solution**
- **High Density Computing Thermal Assessment**
- **IBM Electronic Service Agent™**
- **Optimized Airflow Assessment for Cabling**
- **Scalable Modular Data Center**

Additional services are available, including **hardware warranty upgrades** and factory-installed **Product Customization Services (PCS)**. For more on these services, contact an IBM sales rep.

Conclusion

IBM is the world's most innovative IT company. We've delivered the most powerful supercomputers and server innovations ever, and for fifteen years straight we've led the industry in IT patents⁹. (In fact, we've been issued more than double the number of patents of #2 HP.) We apply our deep heritage and rich business portfolio to develop insights across a variety of industries and create integrated solutions for your more difficult business problems, to help you gain a competitive advantage. In the x86 market, powerful technology from IBM Research has surrounded commodity parts with an unprecedented level of added value. Our goal is simple: to provide you with the best investment value for your technology dollars. To this end, the IBM X-Architecture blueprint—which encompasses hardware, software, and services—provides you with an incredible degree of flexibility, integration, data center simplification, performance, availability, and serviceability. This helps your servers run faster and cooler, use less energy, and be superior virtualization platforms.

These capabilities combine to help reduce your total cost of ownership. Equally importantly, they can help you sleep soundly at night, knowing that your data center is functioning as smoothly, efficiently, and cost-effectively as possible, in a green and sustainable manner. x86 technology has grown into serious mainstream, mission-critical technology, and IBM X-Architecture servers are the best value in the marketplace. Can you afford to be saddled with a weak technology partner? Can you afford *not* to choose IBM?

⁹ <http://ibm.com/press/us/en/pressrelease/23280.wss>.



For More Information

IBM System x Servers

ibm.com/systems/x

IBM BladeCenter Server and options

ibm.com/systems/bladecenter

IBM Rack Configurator

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

IBM ServerProven Program

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

IBM Technical Support

ibm.com/server/support

IBM Configuration and Options Guide

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, Active Energy Manager, AIX, BladeCenter, Calibrated Vecteded Cooling, Chipkill, ClusterProven, Cool Blue, DB2, Electronic Service Agent, eXtended I/O, iDataPlex, Memory ProteXion, Power, POWER6, Predictive Failure Analysis, ServeRAID, ServerGuide, ServerProven, System Storage, System x, Systems Director Active Energy Manager, Tivoli, UpdateXpress, X-Architecture, XceL4v, Xcelerated Memory Technology, and XpandOnDemand are trademarks of IBM Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://ibm.com/legal/copytrade.shtml>.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc.

InfiniBand is a registered trademark of InfiniBand Trade Association.

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

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, and SQL Server are trademarks or registered trademarks of Microsoft Corporation.

TPC Benchmark is a trademark of the Transaction Processing Performance Council.

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

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS

OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

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

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

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

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