How IBM X-Architecture technology can help your business prosper
March 2007

IBM X-Architecture Technology: A Story of Innovation for x86-Based Servers

By Mark T. Chapman
IBM Systems and Technology Group

Executive Overview

Look closely. Innovation ahead.

At IBM®, we rely on deep business experience, renowned research capabilities, and world-class technology to help businesses move forward via innovation. IBM's ultimate goal is to provide innovation that matters to *you*—innovation that helps you solve old problems in new and better ways. However, innovation is only the process; success is the result.

According to IDC[1], nearly 70% of IT server spending in 2010 will be on x86 servers. Once, there was a perception in the industry that these servers were suitable primarily for filling low-end computing niches, such as file-and-print, e-mail and departmental serving. IBM was among the first vendors to envision that x86 servers could take on a much broader role. In 1998, recognizing that x86 servers needed many of the same reliability, serviceability, availability, and performance attributes as mainframes and other advanced servers, IBM introduced x86 servers incorporating X-Architecture® technologies. These game-changing technologies, including x86 industry firsts such as Chipkill™ memory protection, light path diagnostics and Predictive Failure Analysis®, helped elevate x86 servers from their limited roles to their present status as mission-critical, mainstream enterprise servers. The competition has been playing catch-up ever since.

This paper describes the evolution of X-Architecture technology and demonstrates how the melding of industry-standard components with IBM innovation produces servers that offer you a unique set of capabilities with which to effectively manage your business. You'll see why IBM systems equipped with X-Architecture technologies run faster—yet cooler—use less power, 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

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 power/thermal issues and a growing interest in virtualization. IBM has been visionary in recognizing these trends early and providing solutions via X-Architecture technology to propel you into the future:

- **Memory becomes a key to balanced system performance —** As the number of processor cores increase, system resources must keep up to avoid potential bottlenecks. For example, for quad-core processors to achieve maximum performance, they require four times as much memory as single-core processors. 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.
- **The one-application-per-server paradigm will be broken and utilization becomes the key to data center performance —** Even the smallest servers become SMP (symmetrical multiprocessing)-capable via multi-core processors. Many customers have been in a one-application/one-server deployment model with very low levels of utilization. Utilization will become a direct measure of the ROI of your technology investment. Because of this, IBM is investing in system and data center management technologies to help you achieve higher utilization levels. We will continue to work on developing technologies that will optimize how applications operate. Having better control allows you to proactively move workloads to other assets and temporarily assign 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. Would you rather rely on IBM virtualization technology, which evolved from three decades of experience with mission-critical mainframe virtualization—or hardware and

[1] IDC Q306 Server Forecaster.

software vendors who are figuring out virtualization as they go?

- **Power and thermal efficiency is driven by managing the envelope** — IBM continues to invest in developing new tools that allow us to offer servers with better power efficiency and to interact with sensors mounted within servers and elsewhere in the data center. With power reduction comes cooling reduction. Over time, today's power and thermal issues may no longer be practical concerns for the data center.
- **The work of next-generation corporate *desktops* will be driven by a server —** Security concerns will eclipse TCO as the driver of *re*centralization; 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."

X-Architecture Technology Today

The X-Architecture blueprint takes a comprehensive approach—using innovative hardware, software and services—to solving customer challenges today, and it embraces an evolving design approach to address tomorrow's challenges. IBM takes the best of industry-standard technologies—even helping to shape some of them through participation in standards groups—and combines them with IBM innovations to provide the greatest possible flexibility in IBM System x™ and BladeCenter® servers. The fundamental *premise* of X-Architecture technology will not change over time; however, the underlying *technologies* and how they are delivered will constantly evolve to meet the needs of our customers.

The X-Architecture design philosophy targets five key areas.

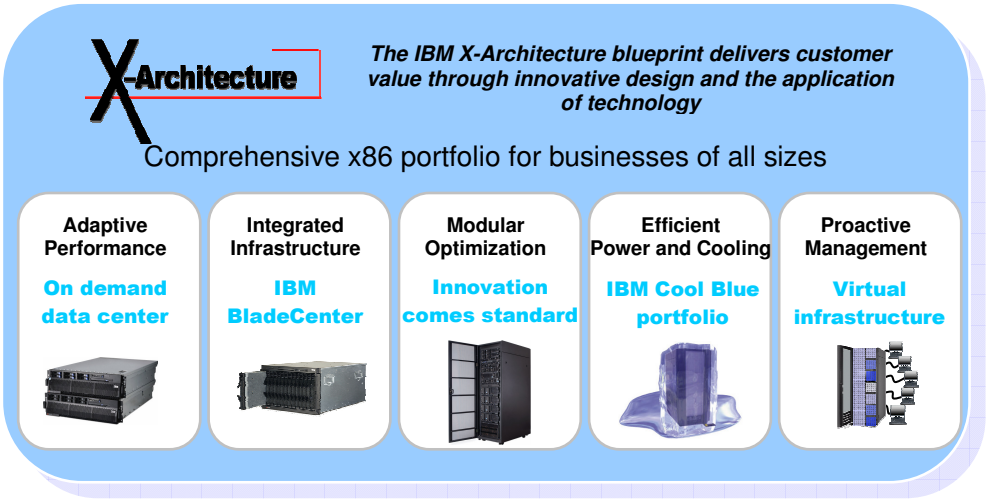

Adaptive Performance

The x86 world is changing. 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. To optimize for this workload, IBM designed scale-up x86 servers with the custom IBM X3 chipset for maximum performance, scalability and reliability. As a result, our state-of-the-art modular building-block

design has led the market for 8-socket-and-up servers for *17 consecutive quarters*, and we hold a **30%** share in 4-socket servers[2]. This demonstrates the strength of our designs.

Today, the line between scale-up and scale-out servers is blurring. 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. Performance becomes 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—so servers and storage resources can be utilized more efficiently.

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 more cost-effective. And it allows the applications to share a dynamically allocated pool of common system resources for improved efficiency.

With all these workloads running on the same server, reliability remains important but *availability* becomes more critical than ever. IBM scalable x86 servers are designed to deliver the highest levels of utilization, flexibility, and availability to the data center, which allows you to adjust your system's performance level as your needs change.

IBM offers servers that are extraordinarily suited for virtualization. For example, using the mainframe-inspired, IBM-designed X3 chipset, the **x3950** with XpandOnDemand™ modular expansion capability is the only *switchless* x86 server from a Tier-1 vendor capable of scaling to **32** sockets. The x3950 is a 3U Intel® Xeon® MP-based 4-socket rack server, with support for up to 64GB of DRAM, 256MB of L4 cache, six high-speed hot-plug adapter slots and six HDDs. If you ever find the need for more resources, simply attach a 3U x3950 MXE chassis to the x3950 using an XpandOnDemand cable. It immediately doubles your processors, memory, adapter slots, HDD bays and L4 cache capacity. Still not enough? Add up to 6 more x3950 MXE chassis to the x3950, for a whopping total of 32 processors/64 cores, 512GB of RAM, 2GB of L4 cache, 48 I/O slots and 48 HDDs in only 24U of rack space.

Alternatively, instead of fully populating the first chassis with four processors and then adding an x3950 MXE chassis with four more, you can choose to populate only *two* of the processor sockets in each chassis, giving those four processors access to 128GB of RAM, 512MB of L4 cache, and **12** high-speed adapter slots. This is the essence of flexibility and optimization.

If you're concerned that all this "external" processor and memory connectivity will hurt performance, don't be. The innovative IBM XceL4v™ Dynamic Server Cache, combined with the custom X3 chipset, boosts performance to the point where the x3950 offers groundbreaking performance in 4-, 8-, and 16-processor benchmarks, whether running Microsoft® Windows® Server or Linux®, IBM DB2® or Microsoft SQL Server®, TPC-C™, SAP, or SAP SD benchmarks.

Again, performance is useless without reliability. **IBM Memory ProteXion**™ technology provides multi*chip* error protection, and works in conjunction with **IBM Chipkill** technology—which provides multi*bit* protection per chip—and standard ECC protection, to provide *three-level* memory correction. Memory ProteXion is standard on x3800, x3850 and x3950 systems.

Two issues that affect virtualization performance are the amount and the speed of memory. Unfortunately, some competitive servers *limit* the amount of memory available to the processors, and some *reduce* the speed of the memory when more than a certain amount is installed.

[2] IDC Q306 Server Tracker.

The current implementation of the Opteron processor supports DRAM speeds of up to 667MHz. Due to the way the Opteron processor is designed, however, when *more than four* DIMM slots are populated per processor, the speed of *all* system memory drops to only 533MHz. The IBM System x3755 avoids this limitation through an IBM patent-pending feature called **Xcelerated Memory Technology**™. Because of this innovation, *all* 64GB of RAM in the x3755 runs at the full 667MHz, using inexpensive industry-standard DIMMs.

A study[3] of relative performance conducted by VeriTest compared the **x3755** server to a similarly equipped HP DL585 G2 server. The results of two different tests clearly show the design advantages of the x3755 with Xcelerated Memory Technology and the **IBM CPU Pass Thru** card[4]. Using identical processor, memory, and software configurations[5], the x3755 with Xcelerated Memory Technology outperformed the DL585 by **10.4%**.

Even more impressively, when the x3755 was configured with only *three* processors and the CPU Pass Thru card in the fourth socket, it outperformed the DL585 with *four* processors[6] by an amazing **27.7%**. Not only does eliminating the fourth processor result in higher performance than the HP can offer, it also affords a **25%** hardware cost savings with a **35%** performance-per-watt advantage, and possibly lower per-CPU software licensing costs.

Integrated Infrastructure

IBM BladeCenter integrates servers, storage and networking to help you reduce complexity, simplify IT management and reduce costs, through a flexible and highly scalable design. The beauty of the BladeCenter architecture is that everything needed for the solution can be housed in a single unit and managed from a single point of control. Because the network switches are all housed inside the chassis, every blade, every switch, every management module, every power supply and every fan can be easily managed from one central source. Instead of having 42 discrete 1U servers in a rack, you can have twice that many servers—using 6 BladeCenter chassis and 84 single-wide blade servers. Yet rather than doubling the number of power, Ethernet, systems management, keyboard, video, and mouse cables—plus KVM switches and PDUs—the number is actually *cut* dramatically.

IBM has held the leading full-year market share in the blade segment for the years 2003-2006.[7] There are four key reasons why customers are selecting IBM over the competition:

- More efficient power and cooling
- Faster I/O throughput and more available ports
- More flexible virtualization
- Better flexibility and investment protection

---

[3] *IBM x3755 AMD Opteron Server CPU Performance and Xcelerated Memory Technology Study*, VeriTest, January 2007.

[4] Although the design of the Opteron allows for **1-** and **3-**processor configurations, it generally does so with a significant performance degradation. The CPU Pass Thru card, which plugs into the second or fourth processor socket, avoids the performance loss, making odd-number processor configurations not only feasible but desirable.

[5] Each server was configured with 4 processors, 32GB of RAM, and Windows Server 2003 EE with SP1, running SPECjbb2005. Results available at *http://veritest.com*.

[6] Each server was configured with 24GB of RAM, and Windows Server 2003 EE with SP1, running SPECjbb2005. Results available at *http://veritest.com*.

[7] IDC Q306 Server Tracker.

Efficient Power and Cooling

Due to the integrated design of BladeCenter's shared power supplies, other devices—such as the internal switches and bridges—may use less power than stand-alone versions. As a result, the aggregate power savings of BladeCenter vs. 1U servers and related external equipment can be as much as **30-40%**[8]. In a packed data center, this can add up to quite an energy savings over the course of a year.

Moreover, instead of having hundreds of tiny fans per rack—using power, subject to failure, and creating ambient noise—BladeCenter requires only *two* **hot-swap/redundant** blowers to cool all the blades and other devices in the BladeCenter chassis. These blowers draw only **60W** *between them*. By contrast, some blade designs require *dozens of non-hot-swappable* fans per chassis, consuming *hundreds* of watts of power and generating lots of noise. Fewer points of failure and less power consumed can mean greater uptime and lower costs. For more information, go to: http://**ibm.com**/systems/x/about/power/bladecenter.html.

Also, you don't have to decide between buying lower-cost servers with nonredundant power supplies versus paying more for redundancy. Likewise, there are only two—optionally four—high-efficiency power supply modules per chassis, providing redundant power for every device in the chassis. Again, this means fewer parts to fail and less power consumed.

Faster I/O Throughput

2006 was the year of the processor. The industry went from single- to dual- to quad-core. IBM introduced several new blades incorporating Intel, AMD, IBM PowerPC® and Cell Broadband Engine™ breakthrough technologies. But without smart system design affording *balanced* performance, processing power can be wasted. With the **BladeCenter Virtual Fabric Architecture** you get the intelligent I/O needed to unleash the power in your processors

IBM is first with full 10Gb Ethernet. Through the **Nortel 10Gb Ethernet Switch Module for IBM BladeCenter**, the BladeCenter H and BladeCenter HT chassis support high-speed 10Gb Ethernet. This means we give you **ten times** the Ethernet performance of HP. In addition we offer the **Cisco Systems 4X InfiniBand Switch Module**. We were first with 4X InfiniBand and we deliver **twice the number of InfiniBand connections per server** than HP. And we have up to **40%**[9] greater bandwidth **per blade** than HP. Even using these faster components, IBM BladeCenter solutions *still use less power and produce less heat than HP's* blade solutions that use slower switches.

In addition to the high-speed switches, all BladeCenter chassis support a number of modular 1Gb switch and bridge modules. 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 *additional* 1Gb or 10Gb ports can be added. Currently, 1Gb switch options include Ethernet, InfiniBand™, Fibre Channel, Myrinet, iSCSI and others from multiple vendors—both IBM and OEM.

When installed in a BladeCenter H high-speed switch bay, the optional **IBM Multi-Switch Interconnect Module** (MSIM) *doubles* the number of ports available to *any* blade server in the chassis. Depending on the blade, this can mean up to **8** or **12 Gigabit Ethernet** or **Fibre Channel** ports per blade.

QLogic Ethernet and **Fibre Channel InfiniBand Bridge** modules, another industry first, provide gateway functions in conjunction with the Cisco 4X InfiniBand Switch for one or more InfiniBand-connected group of BladeCenter H or BladeCenter HT systems to external Gb Ethernet or Fibre

[8] According to comparisons of similarly equipped servers using the IBM System x and BladeCenter Power Configurator.
[9] 20 lanes on BladeCenter HS21, HS21 XM, LS21, and JS21 vs. only 14 lanes on HP BladeSystem half-height blade design.

Channel networks. This further **lowers the cost of ownership** from an acquisition cost and total cost of ownership perspective.

Virtualization Leadership

So why is I/O bandwidth important? There are many reasons. One is virtualization. Virtualization means more applications per server, which puts more demands on the system.

The **BladeCenter H and Cisco VFrame Solution** delivers better virtualization flexibility than HP Virtual Connect does. The IBM/Cisco solution is available today and offers:

- Fabric sharing and I/O consolidation; HP offers no fabric consolidation
- The IBM offering is high-speed, including 10Gb Ethernet and 4X InfiniBand—we also support 1Gb Ethernet, 4Gb Fibre Channel, SAS, and iSCSI; HP's is low-speed—1Gb Ethernet and 4Gb Fibre Channel only
- IBM's solution is based on industry standards and is compatible across all BladeCenter chassis and many OEM switches; HP's design is proprietary and limited to only the c-Class chassis and specific HP switches
- The IBM solution scales to **512** servers; HP's to just **64** c-Class blades and **4** chassis
- IBM's offering is simple to use and virtualizes the server, workload and network connectivity, offering policy-driven automated blade failover; HP's is strictly manual and limited to the network addressing of the blade

Another offering, **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 port addresses used by their server blades via software as an alternative to the addresses that are burned in to the hardware during manufacturing. BladeCenter Address manager is planned for the second-half of 2007.

Flexibility and Investment Protection

BladeCenter sports the largest ecosystem of any blade vendor, which allows us to offer the widest array of IBM and OEM options to our customers. We have made the I/O specifications for BladeCenter available to OEMs. As a result, nearly **100 companies** are already producing solutions for BladeCenter.

To further advance next-generation thinking around blades, IBM formed **Blade.org**. It started as a group of companies interested in driving innovation on the BladeCenter-based platform, and has become a fully functioning organization, supporting openness and collaborative innovation around new products and systems. Blade.org has invested more than $1 billion in the last two years to fund companies developing emerging technologies and solutions to simplify computing with blade servers. In just a year, blade.org has grown to a membership of nearly 100 and includes leaders in blade software and hardware, developers, distribution partners and end users. Collectively, the community has introduced more than 30 Blade.org solutions to the market to help customers choose with confidence from the blade portfolio.

To protect your investment, IBM has announced four chassis since 2002, and all are still being sold. Because of our compatibility across chassis, we've delivered a durable infrastructure that protects your investment and gets your new technologies up and running faster. IBM's flexible, forward-thinking design is part of the reason for BladeCenter's success. By comparison, HP has introduced four *incompatible* chassis and discontinued three of them. This means they have forced their customers to repeatedly start over.

We build innovation and investment protection into every *blade server* as well. The **IBM BladeCenter HS21 XM** is designed for the user who needs the ultimate in processor power, memory capacity, and I/O bandwidth. It supports up to **two *quad-core* Xeon** processors, **32GB**

of memory, and **12 communication ports**—using a **Multi-Switch Interconnect Module** installed in the chassis. It would be difficult to run out of communication ports with the HS21 XM.

And, if your needs run to *scalable* blade servers, there's the **AMD Opteron LS41 for IBM BladeCenter**. It's a standard single-wide (30mm) 2-socket blade server that also supports up to 32GB of DRAM. Should you need more processing power, a *second* blade—an optional 30mm **Multiprocessor Expansion Unit**—can be snapped onto the first in the adjacent blade slot. This *doubles the processor and memory capacity* of the now-60mm-wide blade, to **four** sockets and 64GB of RAM. You don't have to worry about running out of processor and memory headroom with the LS41.

Modular Optimization

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 dual- and quad-core performance, high availability, scalability, power efficiency, and proactive manageability.

For example, the IBM System **x3650**, a 2U server, holds up to 12 DIMMs, for a total of up to 48GB. By comparison, the HP ProLiant DL380 G5 and DL385 G2, as well as the Dell PowerEdge 2950, are limited to only 8 DIMMs. The additional DIMM slots give you the flexibility to use lower-capacity/lower-cost DIMMs for many configurations. The additional memory capacity also contributes to the x3650's leadership results in the **SPECjbb2005** benchmark[10].

For those concerned with the security of their data, the **x3650** 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.

Like the 4-socket x3755, the 2-socket **1U x3455** offers **Xcelerated Memory Technology**. Because of this innovation, *all* 48GB in the x3455 run at the full 667MHz, unlike competitive systems.

Performance is important, but it's only useful when a server is up and running. Consequently, reliability, availability and serviceability are more important than ever. To this end, IBM incorporates a number of innovations as standard features in its System x and BladeCenter servers.

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 panel with an LED for each major component—processor, memory, HDDs, adapter slots, etc. When the servicer opens the system cover, another LED beside the specific component—such as DIMM Slot 12—identifies the failed part. There is no fumbling about, trying to determine which of the 32 DIMMs needs replacement. And light path diagnostics works even when the server is unplugged. Consider the amount of time light path diagnostics could save your technicians over the lifespans of dozens or hundreds of servers. Plus, shorter servicing time can mean less downtime for your servers.

[10] 218,032 SPECjbb2005 business operations per second (SPECjbb2005 bops) and 109,016 SPECjbb2005 bops/JVM. Results as of January 31, 2007, available at http://www.spec.org.

Another innovative feature used in several System x servers is **IBM Xtended 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 your 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. Should you buy a 1U server with two *PCI-X* slots to accommodate your current adapters, or one with two *PCI-E* slots to accommodate your future needs? Or maybe compromise on a server with one of each?

With Xtended I/O, you needn't worry about getting stuck with the wrong configuration. Xtended I/O provides adapter slots on replaceable riser cards. This means you can start out with PCI-X slots today and later replace some or all of them with PCI-E slots as your performance requirements increase. Or start with PCI-E and switch if your needs change. In the case of our Opteron-based 1U and 2U servers—x3455 and x3655—Xtended I/O even includes an **HTX** riser option. HTX adapters offer similar throughput to x16 PCI-E adapters, but with even lower latency.

The snarl of cabling behind most racks is at best inconvenient to work around and at worst an expensive logistical nightmare, requiring the rewiring of servers, PDUs, KVM switches, and other equipment whenever a rack server is added or removed. Even worse, the veil of cables blocks rack airflow and can actually contribute to equipment failure due to overheating. **IBM Advanced Cabling Technology (ACT)** is the solution for reducing behind-the-rack cabling by as much as **87%**. Conventional cabling has bulky KVM cables exiting each server, which then connect to a KVM switch. The cables exiting a series of KVM switches must then be aggregated via additional KVM switches and PDUs, which only increases the number—and cost—of cables, KVM switches and PDUs. Instead, the daisy-chain approach of ACT cabling uses readily available, inexpensive CAT5 cabling to considerably *reduce* the number of cables, KVM switches, and PDUs needed, rather than increasing them. If a server is removed or added, no complicated rewiring is needed. One cable connects the first server in the rack to the next, and so on. Up to 16 servers form a chain; up to 8 chains can connect to one Local Console Manager (LCM); 16 LCMs can connect to one Global Console Manager (GCM). In this manner, up to **2,048 servers** can be centrally managed. Equally importantly, with ACT—unlike some other offerings—everything is done externally via cabling; *no* special adapters are required.

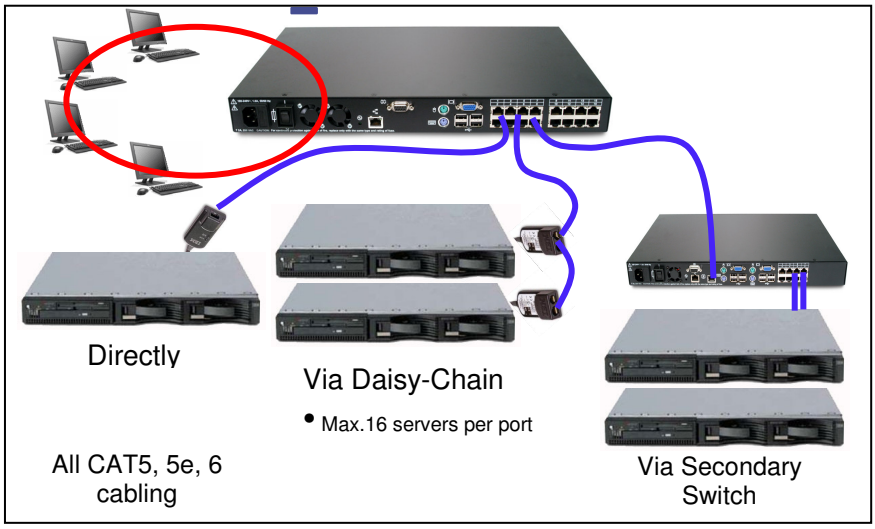

Efficient Power and Cooling

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

According to IDC[11], it would cost almost $400,000 annually to power a data center containing 1,000 x86 servers. At the current industry pace, it could soon be cheaper to build a new data

[11] IDC, *Server Power Consumption Reemerges as a Critical Cost Factor In Data Centers*, Doc #33937, August 2005.

center rather than try to upgrade existing data centers to accommodate increased needs for power and cooling.

A 2006 report[12] by the Robert Frances Group says that ". . . power will be the number one issue for most large company IT executives to address in the next 2-4 years. Ignoring this issue will NOT be an option. Power consideration must be incorporated into data center planning. This includes close cooperation with facilities management, intelligent reduction of unneeded applications, systems, and infrastructure, and re-evaluating system design with an emphasis on power efficiency as a fundamental design criterion."

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 power 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 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.

  To find out how Georgia Tech University uses the **IBM Rear Door Heat eXchanger** to cool racks containing one of the world's most powerful supercomputers, read the press release at: http://**ibm.com**/press/us/en/pressrelease/19231.wss or view the *video* at http://**ibm.com**/systems/x/about/power.html

- **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 provide increased air pressure—angling fans to increase the cooling for specific components, and designing patented *hexagonal ventilation holes* in the system bezel, instead of round ones, because they can be grouped closer together for greater airflow.

  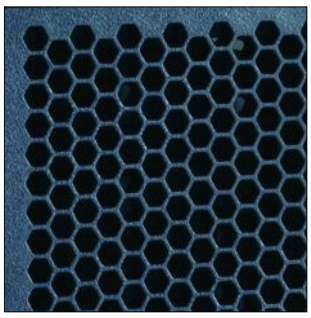

  Another thermal implementation, *isolated zone cooling*, requires only *some* fans in a specific "cooling zone" to switch to full speed in response to higher temperatures in that zone, rather than shifting *all* of them into high gear. Because of this design, the fans use less power and create less ambient noise.

  All this engineering is designed to increase system reliability and availability; but it also enables System x and BladeCenter servers to support more DIMMs and/or more disk drives in some servers than the competition can manage[13]. It also allows some servers to be shallower than most, without overheating. The smaller chassis size leaves more air space behind the servers in a rack. This in turn promotes better rack-level airflow—and thus more effective cooling—as well as simplifying cable management.

- **Energy-efficient power supplies** — The typical power supplies used in the server industry are approximately **65-75%** efficient at converting AC wall current into the DC power used inside a server. This means for every 1,000 watts consumed by the server, perhaps only **700W** are used productively and 300W do nothing more than generate waste heat. 300W equals **1,023** BTUs per hour of hot air that needs to be cooled, wasting even more power.

  By contrast, the power supplies IBM uses in System x servers and BladeCenter chassis are significantly more efficient—up to **91%** efficient in the case of BladeCenter H. This means that for every 1,000 watts of power consumed by the server, you would use **910W** for

[12] *The Rise to Power… of Power: Dealing with the New Data Center Constraint*, Jerald Murphy, Robert Frances Group, 2006.
[13] For example, the 1U **x3455** supports up to **12** DIMMs, vs. only **8** for the HP ProLiant DL145 G2 and only **6** for the Dell PowerEdge SC1435.

processing and waste only 90W generating heat. This helps save you money both on power 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 power. IBM offers a number of blade and System x servers that use as little as **31** to **72** watts per processor. Compare this to the standard Xeon and Opteron parts running at 95 or 103W apiece, respectively. In addition, dual-core and quad-core processors provide as much as double or quadruple the performance at the same power consumption as a single-core version of the same processor. Not all server vendors offer these low-power processors, costing you money on wasted power and cooling.

The innovative and efficient cooling designs of System x and BladeCenter have helped make us more energy efficient than the competition[14]. For example, internal IBM testing reveals that IBM **BladeCenter H** combined with an **LS21** Opteron-based blade consumes up to **24% less** energy than the comparable HP BladeSystem and BL465c blade[15]. Similarly, a **BladeCenter H** and **HS21** Xeon-based blade combo consumes an average of **19% less** than the HP BladeSystem and BL460c blade[16].

IBM BladeCenter & LS21	HP BladeSystem & BL465c	IBM BladeCenter & HS21	HP BladeSystem & BL460c
3 enclosures, 36 servers, 72 processors			
27U	30U	**27U**	30U
5,220W	6,840W	**9,360W**	11,485W
17,732 BTUs	23,324 BTUs	**17,732 BTUs**	31,917 BTUs

IBM PowerExecutive

In order to put control of processor power-saving features at the fingertips of administrators, IBM developed IBM **PowerExecutive**™. PowerExecutive is a powerful software tool, designed to take advantage of new processor features, such as balancing the performance of the system according to available power input. A plug-in for IBM Director—see the *Proactive Management* section—PowerExecutive provides the ability to plan, predict, monitor and cap power consumption based on your System x or BladeCenter hardware configuration. It also allows you to reduce the infrastructure required for redundancy, by using fewer servers with smaller power feeds, and potentially lowering your overall data center support costs. Go to http://demos.dfw.ibm.com/servers/Demo/IBM_Demo_IBM_Director_Power_Executive_Tool-Feb06.html for a demonstration.

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: power 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. Go to http://**ibm.com**/systems/bladecenter/powerconfig to download the Power Configurator.

[14] Power figures based on testing performance on 11/14/2006; the number quoted is the average power difference idle through maximum utilization.
[15] Each server tested with two 2.4GHz Rev F dual-core AMD Opteron processors, 16GB of memory, dual Ethernet controllers, and one 36GB SAS HDD.
[16] Each server tested with two 2.66GHz dual-core Intel Xeon processors, 16GB of memory, dual Ethernet controllers, and one 36GB SAS HDD.

Thermal Diagnostics

IBM Thermal Diagnostics, a planned future plug-in for IBM Director, will provide clients with the additional security capability needed to monitor heat emissions in the data center and determine their root causes—such as blocked airflow or air conditioning failures—before they flare up. Thermal Diagnostics scans data center equipment periodically to collect inventory, performance and temperature metrics. The software then builds a virtual model of the "most-likely scenario" so it can respond automatically to the diagnosis to correct heat-related problems quickly and effectively, or, if necessary—to alert the system administrator of an impending problem.

IBM Global Technology Services

IBM offers a variety of services aimed at helping you to manage and optimize your power and thermal issues. Refer to the *Services* topic, page 14.

Proactive Management

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

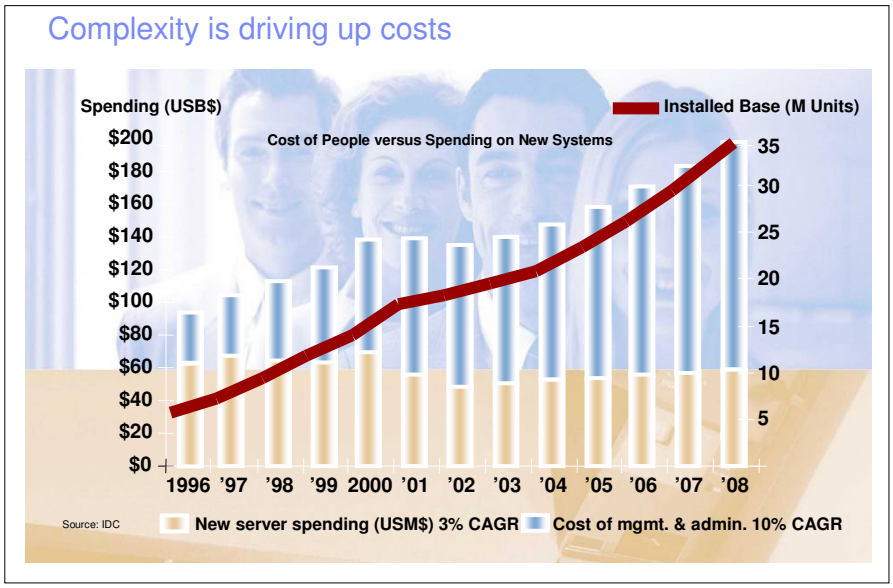

According to an IDC study[17], data center complexity is driving up costs. While new server spending is expected to increase by only 3% per year through 2008, the cost of management and administration is increasing by a whopping **10%** per year.

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

- **IBM Director** for advanced workgroup management is included with all BladeCenter and most System x servers. From a single user interface, IBM Director enables monitoring and event management across a heterogeneous IT environment. From a single access point, users can monitor system resources, take hardware and software inventory, view and respond to events, perform task management, take corrective actions, issue distributed commands and remotely control both servers and storage. IBM Director enables you to customize thresholds and monitor system components—for things like temperature, voltage regulation, etc.—to help maximize uptime. If problems arise, IBM Director can send alerts to administrators via e-mail, cell phone, pager and other means. For more information about IBM Director, agents, extensions and upward integration modules, go to http://**ibm.com**/systems/management/director/about.
- **IBM Remote Deployment Manager (RDM)** supports the deployment, update and retirement of servers, personal computers, workstations and point-of-sale terminals across the computing environment. Based on industry standards, RDM simplifies and automates the deployment, redeployment, configuration management and disposal of x86 systems across

[17] *On-Demand Enterprises and Utility Computing: A Current Market Assessment and Outlook*, IDC #31513, July 2004.

the entire enterprise. It provides high-performance cloning, recovery and reprovisioning solutions, and gives administrators the ability to download an entire system image or a firmware update remotely to any offline node in the network. All of this functionality is provided as an IBM Director module. An optional **extension for IBM Director**, RDM acquires access to all of the tools that IBM Director offers, such as scheduling and event action planning for deployment.

- **IBM UpdateXpress**™, a no-charge Web-based tool, can help reduce your cost of computing by providing a simple yet effective way to update server firmware and the firmware of supported options within supported servers. With UpdateXpress, administrators no longer have to spend hours visiting each individual server that requires updating. Now they can update the server or client system firmware from anywhere on the network. Download UpdateXpress from http://**ibm.com**/systems/management/xpress.html.
- **IBM Virtualization Manager** is an optional **extension for IBM Director**. Virtualization Manager allows you to manage physical and virtual machines from a single console. Combining IBM Director and Virtualization Manager, you can manage both VMware ESX Server and Microsoft Virtual Server environments. Virtualization Manager also integrates VMware VirtualCenter and IBM Director for advanced virtual machine management.
- **IBM PowerExecutive**, described previously, is a no-cost **extension for IBM Director** utility that allows the administrator to plan, predict, monitor, and even limit power consumption by System x and BladeCenter servers.
- **IBM Dynamic System Analysis (DSA)** collects and analyzes system information to aid in diagnosing system problems. In addition, DSA creates a merged log that allows users to easily identify cause-and-effect relationships from different log sources in the system. DSA also compares device driver and firmware on the system to the versions available on an UpdateXpress CD, providing a summary of the differences.
- Extensive **Predictive Failure Analysis** (**PFA**) support provides up to 48 hours 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. On specific servers, PFA support is enabled for processors, memory, hard disk drives, power supplies, fans, voltage regulator modules, and the x3950's XceL4v Dynamic Server Cache. By comparison, many competitive systems offer PFA *only* for memory, HDDs, and possibly processors.
- Most System x servers include an integrated **Baseboard Management Controller** (BMC)[18] which alerts IBM Director to anomalous environmental factors, such as voltage and thermal conditions—even if the server has failed. Some servers have a miniBMC, which provides a subset of these capabilities.
- A "slotless" **IBM Remote Supervisor Adapter II SlimLine**—standard in some models; optional in others—provides additional systems management capabilities for most System x Servers, 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; and other features.

A new tool, planned for late-2007 release, is the **IBM Continuous Availability Manager (CAM)**, an **extension for IBM Director**. It 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. Other features include local/remote virtual machine (VM) restart, local/remote VM checkpoint restart, VM migration, and container/VM record and replay. Its "hypervisor-neutral" design allows it to work equally with Xen, VMware, Microsoft Virtual Server, or none at all. CAM is transparent to the OS, requires no internal agents and supports multiple open interfaces.

[18] The BMC is IPMI 2.0 or 1.5-compliant (system-specific); the miniBMC is ASF 2.0-compliant.

Storage

IBM offers a number of innovative storage solutions, both inside and outside the server, to enhance data security and availability, increase performance and improve storage TCO:

- **IBM System Storage™ DS3000 series** — A powerful and affordable new storage family with integrated management software designed to meet SMB and departmental needs. The DS3000 series is a perfect antidote to SMBs' largest headache: rapid and uncontrolled proliferation of data. The **DS3200** and **DS3400** each offer single-controller and dual-controller models to help maintain data-availability. The DS3200 is direct-attached storage with support for clustering and a 3Gbps SAS connection to the host. The DS3400 can be direct-attached or used in a SAN configuration via 4Gbps Fibre Channel connectivity. In addition, the DS3000 expands to 48 drives—via 3 optional **EXP3000** expansion units—and **14.4TB** of storage. DS3000 **a**ttaches to System x, BladeCenter, *and* select **HP and Dell** servers.

  Dell/EMC's entry Fibre Channel product, AX150, and HP's MSA 1000 and MSA 1500 provide only *half* the throughput, at only 2Gbps. Even Dell's newer SAS MD3000 cannot support more than 15 internal disks[19]—**33 *fewer*** drives than is supported in the DS3000 series.
- **RAID-1E/1E0/5EE** — IBM developed additional RAID levels—available in select **IBM ServeRAID™** controllers—to add functionality and data security beyond the standard RAID levels. These include **RAID-1E** enhanced mirroring using odd numbers of drives, **RAID-1E0** enhanced mirroring combined with data striping, and **RAID-5EE** enhanced RAID-5 with hot-spare drive support.
- **Simple-swap SATA drives** — Typically, servers offer a choice of HDD implementations: either hot-swap drives or fixed drives. Hot-swap drives offer convenience and higher availability due to shorter servicing times, but at a higher cost. Fixed drive models are less expensive, but at the cost of inconvenience and lower availability. IBM offers an alternative in some models: simple-swap drives. They offer the convenience of hot-swap drives, by using a drive tray that slides in and out of the front of the server without requiring tools and at a lower cost than hot-swap drives. Like fixed drives, the server must first be powered off.
- **Storage virtualization** — The increased spread of Fibre Channel™ and iSCSI SANs will drive the need for storage virtualization. IBM System Storage products can help transform the economics of enterprise storage by enabling users to simplify their infrastructure, protect their data and efficiently manage information. For organizations interested in increasing flexibility while reducing the cost of storage, TotalStorage virtualization capabilities allow you to virtualize disk, file and tape. Our virtualization solutions for System x and BladeCenter include the **TotalStorage™ SAN Volume Controller**, **TotalStorage SAN File System**, **IBM Tivoli® Intelligent Orchestrator** and **IBM Tivoli Provisioning Manager**.

Services

IBM offers a number of innovative services to help you optimize and maintain your BladeCenter and System x solutions and overall data center:

- IBM **Electronic Service Agent™** is a "call home" feature that allows System x servers to automatically report hardware problems to IBM support, which can even dispatch onsite service if necessary[20]. 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. Go to http://**ibm.com**/support/electronic to download[21] this no-charge tool. The link can be found on the left side of the page, under IBM Electronic Services.
- **Data Center Consolidation and Relocation** — Provides a consistent, repeatable, phased

[19] *http://www.dell.com/content/products/productdetails.aspx/pvaul_md3000?c=us&l=en&s=bsd&cs=04*.
[20] For those customers entitled to onsite support under the terms of their warranty or an IBM Maintenance Agreement. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.
[21] IBM Service Agent can run either as an extension for IBM Director or stand-alone.

management approach to help you implement a data center relocation that crosses international boundaries.

- **High Density Thermal Assessment** — Helps you identify and resolve heat-related issues within your data centers and provides a set of scenarios for future additions.
- **Scalable Modular Data Center** — Enables you to create and rapidly deploy a data center using predefined solutions. The capability can be installed in nearly any working environment and provides ready power, cooling, security and monitoring for the data center environment.
- **High Density Readiness Assessment** — Helps you gauge your capacity to support high-density IT infrastructure components in their data center facilities, identifies potential gaps that could jeopardize continuous operations, and provides a foundation for the development of remediation measures.
- **High Density Integrated Rack Solution** — Helps you design, deploy and manage resilient, flexible racking solutions to take advantage of new technologies within new and existing environments.
- Additional services are available, including **hardware warranty upgrades** and factory-installed **Product Customization Services** (PCS), such as asset tagging, hardware integration, software imaging and operating systems personalization.

Conclusion

Last year, IBM consultants interviewed more than 450 CEOs around the world. **Eighty** percent said their primary objectives have shifted away from cost-cutting to **revenue growth driven by innovation**. In striving for growth, most say their companies are neither responsive enough to changing business conditions nor agile enough to pursue new market opportunities. At IBM, we're passionate about removing these roadblocks for our customers. We're equally passionate about innovating products, systems, processes, and integrated solutions to make you more responsive, and, in turn, more differentiated and competitive in your own markets.

IBM is the world's most innovative technology company. We've delivered some of the most powerful supercomputers and server innovations ever, and we continue to lead the industry in IT patents. 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 enduring value. Our goal is simple: to provide you with the best investment value for your technology dollars. To this end, IBM **X-Architecture technology**—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 power, 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.

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?

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

Produced in the USA
March 2007

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, BladeCenter, Calibrated Vectored Cooling, Chipkill, ClusterProven, Cool Blue, DB2, Electronic Service Agent, Memory ProteXion, PowerExecutive, PowerPC, Predictive Failure Analysis, ServeRAID, ServerProven, System Storage, System x, Tivoli, UpdateXpress, TotalStorage, X-Architecture, XceL4v, Xcelerated Memory Technology, XpandOnDemand and Xtended I/O are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see http://**ibm.com**/legal/copytrade.shtml.

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

InfiniBand is a 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.

UNIX is a registered trademark in the U.S. and/or other countries licensed exclusively through The Open Group.

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.