How IBM is working to save you money on power and heat-related issues
September 2007

IBM System x and BladeCenter Power and Thermal Advantages

By Mark T. Chapman
IBM Systems and Technology Group

Executive Overview

There is no question that running a data center is expensive. Hardware, software, floor space, people—they all add up. And every day the demand grows for more speed and greater capacity. Unfortunately, faster processors, larger memory configurations, more disk drives and more I/O adapters add up to more power needed and therefore increased electrical cost.

Then there's the additional cost to remove all the heat produced by the hardware. Heat translates into a shorter lifespan for hardware. The more excess heat is present, the more damage can be incurred by hardware. Thus you either pay to keep the hardware cool, or you pay to replace heat-damaged components. And rising utility rates make the solutions ever more costly.

The goal of harried IT directors everywhere is to reduce the demand for electricity and keep servers and the data center cool. Doing so reduces costs and increases hardware reliability. This paper describes how IBM® products and services aid IT directors in reaching their goal.

What Are Your Options?

Competing pressures in the data center compound the thermal and power issues. The trend toward server consolidation means replacing stand-alone tower servers and larger rack servers with rack-dense 1U and blade servers. This saves floor space and centralizes server administration, but it also increases the heat per cubic inch of space. The more densely the hardware is packed together, the harder it is to remove the heat from the servers. Once expelled by the servers, there is the additional problem of how to remove all that heat from the data center.

This leads to a number of problems that IT directors have to solve:

- How do I get enough power into the data center to keep up with rising demand?
- How do I effectively handle air movement?
- How do I keep everything cool?
- How do I make the best use of the space I have?

Already, many data centers have exceeded their ability to power and cool full racks of equipment. This results in many racks being left partially empty, which wastes valuable floor space. So what can you do to increase your power and thermal efficiency and reduce costs?

There are various approaches that can be taken and IBM is attacking these problems at three different levels:

- Server level
- Rack level
- Data center level

The following topics will illustrate IBM's efforts in each of these areas.

Server-Level Solutions

IBM attacks power and thermal issues at the server level in multiple ways:

Low-Power Components

Replacing high-powered components with lower-power equivalents, allows us to reduce the overall power demand for a server. Specific examples include:

- **Replacing 3.5-inch hard disk drives (HDDs) with 2.5-inch drives —** At the same capacity, 2.5-inch drives consume up to **40%** less power than their 3.5-inch counterparts. Servers that use 2.5-inch drives currently include the IBM System x3250, x3550, x3650, x3850 and x3950 rack-optimized servers, as well as the HS20, HS21, HS21 XM, JS20, JS21, LS20, LS21, and LS41 blade servers for IBM BladeCenter®. Other IBM System x™ servers will migrate to 2.5-inch drives over time.

- **Using low-voltage dual-core processors —** Both Intel® and AMD offer low-voltage versions of some of their processors. They run at the same clock rates as their higher-voltage cousins, but consume less power. In addition, dual-core processors provide as much as double the performance at the same power consumption as a single-core version of the same processor. IBM servers that use low-voltage dual-core processors include many blade servers and System x servers. All of the following servers offer models that consume fewer than 75W per processor:
  – Some models of the HS20 utilize **ultra-low power** Intel Xeon® processors. These **dual-core** processors consume only **31W** of power (**15.5W** per core). They use **70%** less power than the previous-generation 103W Xeon processors, while also delivering excellent price/performance.
  – Some HS21/HS21 XM models employ **35W** (**17.5W** per core) or **40W** (**20W** per core) **dual-core** Xeon processors. They use **63%** and **58%** less power, respectively, than the standard 95W Xeon processors.
  – Many **dual-core** tower, rack, and blade server models offer **65W** Xeon processors (**32.5W** per core), rather than 95W processors (**31%** less power usage).
  – *Some LS20/LS21/LS41* models use **68W dual-core** AMD Opteron processors (**34W** per core), instead of the 95W variety (a **28%** reduction).
  – The *JS20* uses **72W** IBM PowerPC® 970FX™ **dual-core** processors (**36W** per core), which consume almost **37%** less power than 95W Xeon or Opteron processors.

- **Using power-efficient quad-core processors —** Quad-core processors provide up to double the performance of dual-core processors, at a lower *per-core* power usage. Many IBM tower, rack, and blade servers currently offer quad-core processors.
  – The **80W quad-core** Xeon processors offered in many System x and HS21/HS21 XM servers consume only **20W** of power per core. Not only does one 80W *quad*-core processor potentially offer up to **90%**[1] more performance than a 95W *dual*-core Xeon processor (of the same clock rate), it also uses **16%** less power in doing so.
  – Some HS21/HS21 XM/x3550/x3650 servers use **50W quad-core** Xeon processors. These consume only **12.5W** per core, for an even better performance-per-watt ratio.

Energy-Efficient Components

When designing a server, a vendor has various decisions to make. Among these decisions is how energy-efficient the components should be. The more efficient the component, the less power is wasted. IBM strives to make its servers as energy efficient as possible by:

- **Using energy-efficient power supplies —** The typical power supplies used in the server industry are approximately **65-75%** efficient at converting between AC wall current and the DC power used inside the server. This means that 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 of hot air that needs to be cooled for no benefit.

  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. This means that for every 1,000 watts of power consumed by the server, you would use 910W for processing and waste only 90W generating heat. This could save you money both on power consumption up front and on cooling at the back end.

[1] Assumes heavily threaded applications. Lightly threaded and single-threaded applications may see little or no performance improvement.

  - BladeCenter requires only **two hot-swappable redundant blower modules** to cool all the blades and other devices in the BladeCenter chassis. (These modules draw only 100W between them.) By contrast, some other blade designs require *dozens of non-hot-swappable* fans per chassis, consuming hundreds of watts of power. Fewer points of failure and less power consumed could mean lower costs and greater uptime.

- **Using power-saving processors —** IBM employs Xeon processors that incorporate a new feature called *Demand Based Switching*[2] *(DBS) with Enhanced Intel Speedstep® Technology*. DBS is included in all dual- and quad-core Xeon processor-based System x and BladeCenter servers.

  There are often situations where a processor is not fully utilized. This provides the opportunity—if the workload allows it—to run the processor at reduced speed and consequently consume less power. Originally developed by Intel for Xeon processors, and then further optimized by IBM, the DBS feature utilizes the operating system to monitor processor utilization based on the applications running on the server. The system can automatically change to a lower power state (frequency and voltage) when less processing power is needed. For example, an e-mail server may run at capacity during business hours, yet be idle during the evenings and on weekends. Or a server might have surplus capacity that has not yet been tapped.

  DBS provides the ability to dynamically change from peak performance to cost-saving mode automatically. Based on IBM's recent internal WebBench Power measurements, DBS reduced average system power consumption by up to **24%** with the processor operating at approximately 45% utilization. Such reductions in power consumption could reduce Total Cost of Ownership by up to $200[3] per system installed.

  AMD provides a similar capability (called *PowerNow!™ technology with Optimized Power Management*) in its Opteron processors. PowerNow! technology can reduce CPU power at idle by as much as **75%**[4].

<u>Efficient Use of Components</u>

Because dual-core (DC) processors can do as much as twice the work of single-core (SC) processors at the same clock rate, they are an efficient way to increase processing throughput without increasing power consumption and heat output. A server equipped with *one DC* processor may provide approximately the same overall throughput as a server equipped with *two SC* processors running at the same clock rate; however the dual-core processor will use half the power and produce half the heat of the two single-core processors. Similarly, *one* server with *two DC* processors will be able to run approximately the same workload (software permitting) as *two* servers with *two SC* processors apiece, while consuming far less power and producing far less heat than two servers would. By the same token, *quad*-core processors offer up to double the throughput of dual-core processors, using fewer watts-per core in the process.

Dual-core Intel and AMD processors are currently available across the IBM System x and BladeCenter product lines. IBM servers currently offer quad-core Intel processors, with quad-core AMD processor-based models coming soon.

---

[2] DBS is operating system-dependent. To take advantage of DBS, please make sure your operating system supports this feature.

[3] Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit www.intel.com/performance/resources/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104. System based on Intel Server Board SE7520JR2 board, 4 GB DDR2-400 memory. 24% power savings at approximately 46% CPU utilization with DBS off running WebBench. Cost assumptions were made for the systems tested. "Power in" costs were assumed at $0.10/KWh, and "cooling costs" were assumed to be roughly 2X the power-in costs. Customer results may vary depending on the hardware and software configuration and the customer's environment (e.g., climate differences and utility costs).

[4] When the server is running an operating system that supports this feature.

Another way in which to reduce power requirements is to utilize iSCSI or Fibre Channel Storage Area Networks (SANs) or Network-Attached Storage (NAS). By consolidating hard disk drives in one or more central storage units, shared among many servers as needed, the need for hard drives in the individual servers is reduced or eliminated. For example, rather than adding a second HDD to a hundred individual servers for redundancy, a SAN array containing perhaps 28 HDDs may be able to service those same hundred servers for less than 30% of the power cost, and with a much higher utilization percentage of those drives. (In other words, with fewer wasted resources—not to mention 72 fewer drives to potentially fail.) In addition, although the individual servers may contain only 3.5-inch HDD bays, the SAN may use lower-power 2.5-inch drives, for additional power savings.

Another option, for blade server users, is solid-state storage. HS21 XM blade servers offer the ability to use either one or two internal 15.8GB solid-state drives. These drives consume only **1W** apiece, which not only decreases power draw and heat output by up to 95% (vs. 3.5-inch drives), but these drives also offer up to three times the MTBF rate[5] of HDDs.

Efficient Usage of Server Resources

Typical software is unable to keep a processor busy most of the time. In fact, the average utilization of x86 processors (Intel and AMD) is on the order of only **15%-40%**. This means that much of the energy used by the processor is wasted while the processor is idling. As a result, in order to keep up with the demands of your users, you find yourself constantly adding new servers. Next time, perhaps, when you have additional workloads to run, instead of buying another server (which will consume more power and generate more heat and will likely run just as inefficiently as the others), why not try **virtualization** techniques? IBM supports tools such as VMware ESX Server, Microsoft® Virtual Server and IBM Virtualization Engine™ that allow you to partition processors and other server resources (local and remote) so that multiple operating systems and multiple application sets can be running concurrently and independently on the same server.

If four different software stacks, for example, are each assigned 20% of the processor resources, you can achieve up to 80% utilization of the processor cycles (instead of 15%-40%), with headroom to spare. If one or more of those stacks later requires additional resources, or if you need to add another stack, you have the available cycles. Using dual-core processors gives you even more flexibility in this regard, providing twice as many physical cores that can be partitioned into logical processors. This is an excellent way to make efficient use of your processing power, even with single-threaded applications.

More Efficient Designs for Server Cooling

Part of the reason many servers retain so much heat is inefficient design. Providing effective cooling takes more than simply sticking some fans under the covers of a server. It requires a holistic design that takes into account the size and shape of the cabinet, the positioning of components, the location of cables, the types of fans and how they are angled, the airflow pathways—even the shape of the air holes in the cabinet[6]. (*Figure 1.*) Only by designing the server around the cooling needs of the components can an effective design result. This was the thinking behind IBM Calibrated Vectored Cooling™.

Figure 1. *Hexagonal air-intake holes*

[5] According to IBM testing, solid state drives provide up to a 3,000,000 hour Mean Time Between Failure rate, vs. up to 1,000,000 hours for SAS HDDs.
[6] Hexagonal holes can be grouped more densely than round holes, providing greater airflow through the system cover.

The evolution of Calibrated Vectored Cooling brought together engineers with experience in multiple disciplines, including some with backgrounds in advanced cooling design and others with experience in chip module packaging. They collaborated with resident IBM thermal experts to find new approaches to cooling Intel-architecture systems. Some solutions were simple: use tandem counterrotating fans (which offer increased air pressure compared to a single fan of the same size) and angle some of the fans to change the airflow direction and increase the cooling to specific parts.

Other solutions were more sophisticated: add cooling fins to the *underside* of the heat sink (in BladeCenter), and employ advanced heat sink technologies, such as vapor chamber (also called planar heat pipe). Another thermal implementation, isolated zone cooling, requires only fans in a specific “cooling zone” to switch to full speed in response to higher thermal requirements in that zone, rather than shifting *all* fans in the server into high gear. Not only does this reduce ambient noise (often a problem in data centers), it also reduces wear and tear on the fans that were *not* sped up.

It was this design methodology that allows 1U IBM servers, such as the x3455 support up to **twelve** DIMM slots, despite the extra heat generated by the additional DIMMs. Competing servers[7], without Calibrated Vectored Cooling, are limited to only **eight** or even **six** DIMM slots in similar-sized units. Likewise, the 2U x3650, can support up to **12** DIMMs, and the 2U x3655 up to **16** DIMMs, vs. only **eight** for some competitive systems[8]. Similar techniques also allow the x3250 server to be smaller than most 1U servers, at only **22 inches (55.9 cm)** deep, 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. All System x, xSeries, and BladeCenter blade servers utilize Calibrated Vectored Cooling.

Also contributing to greater behind-the-rack airflow is IBM’s optional **Advanced Cabling Technology** (ACT). ACT allows you to daisy-chain together racks of servers and eliminate up to **85%** of KVM (keyboard, video and mouse) cabling, along with KVM switches and some power distribution units (PDUs). ACT reduces the tangle of cables, simplifying cable management, and can considerably lower cabling costs. (*Figure 2.*) ACT even allows you to incorporate standard KVM switches, if needed.

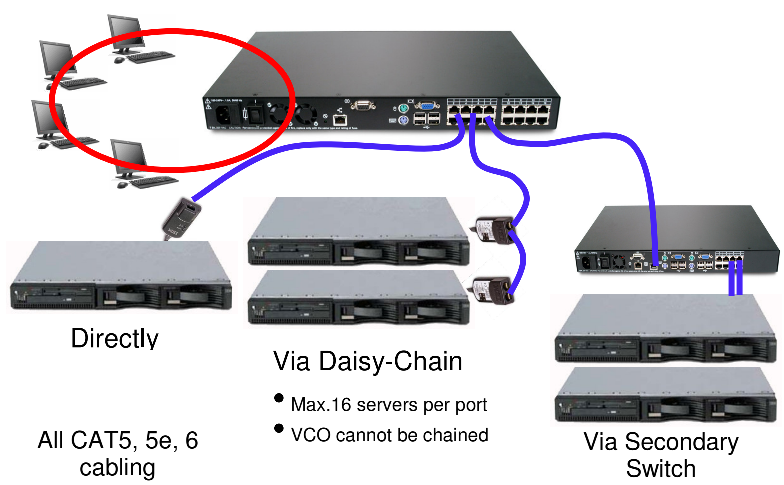

Figure 2. *ACT daisy-chain cabling, interconnecting up to 2,048 servers*

[7] HP ProLiant DL145 G2 (8 DIMMs); Dell PowerEdge SC1435 (6 DIMMs).

[8] HP ProLiant DL380 G5 and DL385 G2; Dell PowerEdge 2950.

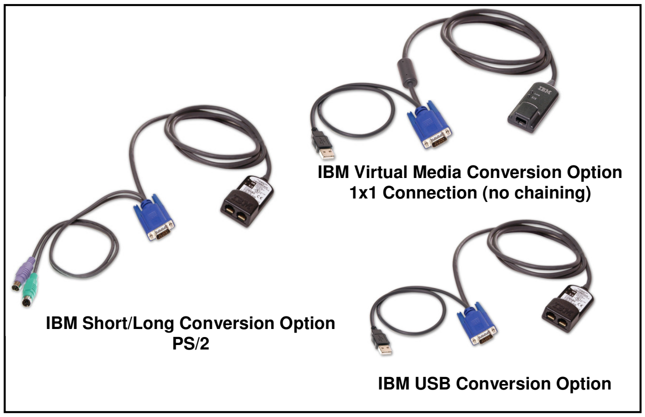

Fewer KVM switches and PDUs can also reduce the data center power draw and cooling demands. ACT is compatible with all System x, xSeries, and BladeCenter servers and requires no internal adapters. *Figure 3* illustrates the various connection options.

Figure 3. *ACT connectors*

Just to keep your servers working efficiently, IBM equips many systems with extensive Predictive Failure Analysis® (**PFA**) support. PFA provides early warning that a component[9], such as a memory or hard disk drive is *about* to fail. This way you can replace the component *before* it fails and compromises the cooling efficiency of a server.

Rack-Level Solutions

IBM doesn't just address power and thermal issues at the server level. Looking at the issue from a rack perspective leads to additional means of combating heat and power problems:

Software to Manage Power and Thermal Issues

In order to put control of processor power-saving features at the fingertips of administrators, IBM developed IBM **PowerExecutive**™. IBM 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.

IBM PowerExecutive provides the ability to plan, predict, 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. It does this, for example, by inventorying all components at the blade level, then adding up the power draw for each blade and tracking that usage. In failure mode, IBM PowerExecutive (via the BladeCenter Management Module) might request that certain blades in each domain throttle down to reduce power consumption. Version 2.0, allows the administrator to impose hard limits on power usage. Go to http://demos.dfw.ibm.com/servers/Demo/IBM_Demo_IBM_Director_PowerExecutive_Tool-Jul06.html for a demonstration of IBM PowerExecutive.

Hardware Consolidation

Typical racks contain many power-consuming devices besides just servers and storage. There may be KVM switches, Ethernet switches, Fibre Channel switches, Myrinet and other high-speed communication switches, plus hundreds of power, KVM and communications cables to link everything together. In addition, each rack-optimized server contains components that consume power, including floppy and CD-ROM drives, systems management adapters, power supplies, fans, and so on.

By consolidating up to four communications switch modules, four power supply modules, two blower modules, two management modules, a CD-ROM drive and a floppy drive into one BladeCenter chassis (containing 14 blade servers), IBM was able to remove *more than a hundred* components from the individual servers and racks and replace them with a few centralized components per BladeCenter chassis. This offers a number of advantages: lower overall power

[9] Depending on the system, PFA can monitor processors, memory, hard disk drives, voltage regulator modules, power supplies, and fans/blowers.

usage, lower heat output, fewer potential points of failure, simplified server management, and extensive redundancy and hot-swappability.

According to IBM internal testing, the BladeCenter design can reduce power utilization by as much as **20%-40%** compared to an equivalent number of 1U servers, and greatly improve air flow behind the rack. (Consolidation of hardware—along with using smaller, lighter hardware, such as 2.5-inch drives—also addresses the nontrivial issue of how much weight a data center's floor can support.)

Room-Level Solutions

Once you have minimized power and heat issues at the server and the rack levels, the next step is to remove the excess heat from the data center as efficiently as possible:

Effective Data Center Heat Removal

With the difficulty of managing airflow characteristics and thermal load issues in many data centers, IBM developed the ***IBM Rear Door Heat eXchanger***. It attaches to the back of a 42U IBM enterprise rack cabinet to provide high-efficiency water cooling right at the rack and help reduce the strain on existing computer room air conditioning (CRAC). The Rear Door heat exchanger is ideal for "hot spot" management and super-high-density applications where the data center cooling solution may be inadequate. It can remove up to **50,000** BTUs of heat (**15KVa**) per rack.

Typically, rack servers will suck cool air in through the front of the rack and blow hot air out the back. The heated air will rise and (hopefully) be captured by the CRAC and cooled before being returned to the room to begin the process again. Unfortunately, this system is rarely as effective as desired. Inevitably, some of the hot air circulates toward the front of the rack where it is sucked back into the upper servers in the rack. As a result, the cooling effect is diminished. (*Figure 3.*)

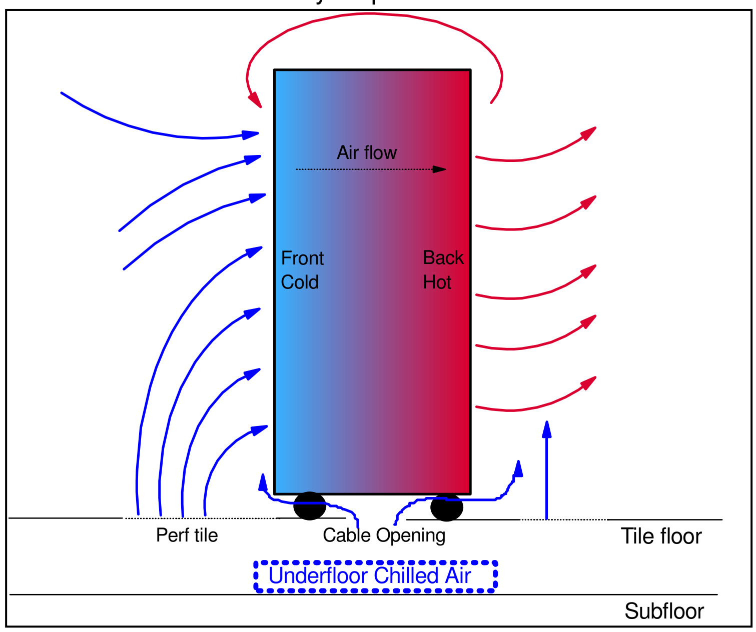

Figure 3. *Typical data center rack*

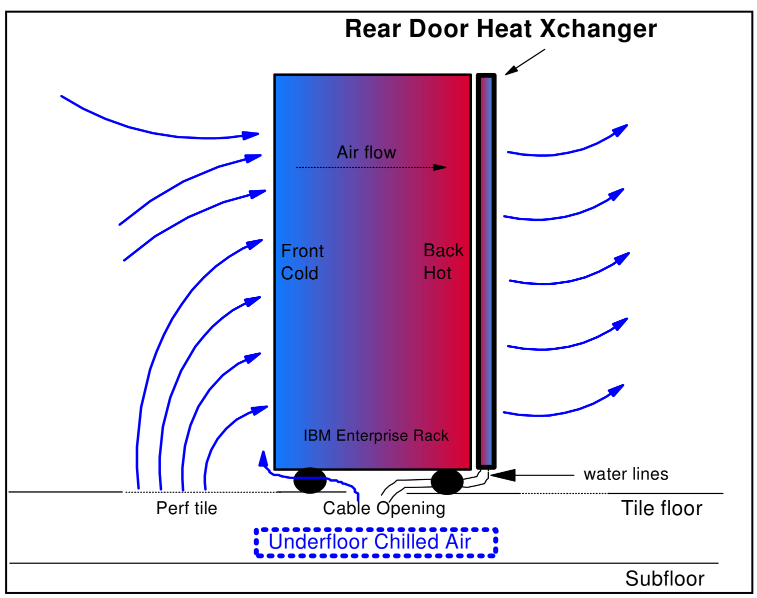

Figure 4. *Data center rack using an IBM Rear Door Heat eXchanger*

By contrast, with the *5-inch deep* Rear Door Heat eXchanger in place, when hot air exits the rear of the servers, the heat is absorbed by the chilled water circulating within the door. (The door requires no electricity and can be opened easily to allow servicing of the devices in the rack.) The water carries the heat out of the data center through the plumbing in the floor, using existing AC chilled water lines. Then the water is chilled again and returns to the Rear Door Heat eXchanger to continue the cycle. (*Figure 4.*)

This way, little if any hot air escapes into the room to be sucked into server air intakes, offering the potential for much cooler server operating temperatures. The energy cost of this solution may be significantly less than a chilled-water AC unit. (Savings may be even greater compared against a DX refrigeration unit.) See **Appendix A** for photos of the unit.

Power Consulting Services

Access to the BladeCenter Infrastructure Support Center (BISC) is a free service for BladeCenter customers. It allows you to contact IBM power and thermal engineers directly. The intent of the service is to help you to understand, optimize and manage your power and thermal load issues, and to think in environmental terms, as well as in system terms. For more information about the BISC, contact your IBM representative.

Side-By-Side Comparisons

So what does all this mean for you? Are all these purported advantages of IBM servers versus the competition merely marketing fluff? Not at all. In fact, we're willing to put our products up against the competition. *Table 1* shows the results of using the Dell (*dell.com*), HP (*hp.com*) and IBM tools provided for calculating power requirements of similarly configured[10] Xeon processor-based blade solutions:

	Typical 1U dual-core Xeon server (average of Dell, HP and IBM 1U offerings)	**HP 1U DL360 G5** (Dual-core Xeon processors)	**Dell PowerEdge 1950** (Dual-core Xeon processors)	**x3550** (Dual-core Xeon processors)	**BladeCenter HS21** (Dual-core Xeon processors)
Number of processors	72	72	72	72	**84**
Rack U-space required	36U	36U	36U	36U	**21U**
Power requirement (watts)	13,752W	15,264W	14,904W	**11,088W**	**9864W**
Heat output (BTUs)	46,878	52,050	50,822	**37,764**	**33,636**
Power/heat savings vs. HP and Dell	N/A	N/A	N/A	**27.3% (HP) 25.6% (Dell)**	**35% (HP) 34% (Dell)**

Table 1. *Power and thermal comparison of Xeon processor-based blade offerings from HP, Dell and IBM*

These results clearly show that while the blade offerings from each of the three companies are more electrically and thermally efficient than an average 1U rack-optimized server, IBM's

[10] All servers configured with **two dual-core 3.0GHz Xeon 5160** processors, **8GB** of memory, and **two Ethernet** ports. Calculations based on running at maximum utilization.

BladeCenter is smaller, uses far less power and produces considerably less heat than the blade offerings from Dell and HP. Similar results can be seen in a comparison of Opteron processor-based blade offerings[11] from IBM and HP. (Dell currently does not offer any Opteron-based dual-processor servers.)

	Typical 1U dual-core Xeon server (average of Dell, HP and IBM 1U offerings)	**HP DL585** (Dual-core Opteron processors)	**x3755** (Dual-core Opteron processors)	**BladeCenter LS41** (Dual-core Opteron processors)
Number of processors	80	80	80	**112**
Rack U-space required	40U	40U	40U	**36U**
Power requirement (watts)	13,752W	11,430W	**7,660W**	**12,148W**
Heat output (BTUs)	46,878	58,170	**26,130**	**41,428**
Power/heat savings vs. HP	N/A	N/A	**33.0%**	**12%**

Table 2. *Power and thermal comparison of Opteron processor-based blade offerings from HP and IBM*

Summary

IBM combines innovative power-saving server designs with energy-efficient and low-power components, hardware consolidation offerings, innovative cooling techniques, virtualization, and intelligent power management software to give you the tools you need to reduce power and cooling requirements in the data center. Added up, IBM helps you simplify your power and thermal management and reduce costs.

[11] All servers configured with **two dual-core 2.0GHz Opteron 8212 HE** processors, **4GB** of memory, and **two Ethernet** ports. Calculations based on running at maximum utilization.

Appendix A.

Figure 5. *IBM Rear Door Heat eXchanger with door closed (L.) and open (R.)*

Additional Information

For more information on IBM System x and BladeCenter directions, products and services, visit our Web site at:

- http://**ibm.com**/systems/x/tower/index.html for more on **x3105**, **x3200**, **x3400**, **x3500**, and **x3800** tower servers
- http://**ibm.com**/systems/x/rack/index.html for more on **x3250, x3455, x3550, x3650, x3655, x3755,** and **x3850** rack-optimized servers
- http://**ibm.com**/systems/x/scalable/index.html for more on **x3950** scalable enterprise servers
- http://**ibm.com**/systems/bladecenter/products for more on **BladeCenter HS20/HS21/HS21 XM/JS20/JS21/LS20/LS21/LS41** blade servers
- http://**ibm.com**/systems/bladecenter/chassis/index.html for more on the **BladeCenter chassis** and options
- http://**ibm.com**/servers/eserver/xseries/storage/pdf/IBM_Rear_Door_Heat_Exchanger_FINAL.pdf for more on the IBM **Rear Door Heat eXchanger**
- http://**ibm.com**/servers/eserver/xseries/storage/rack.html for more on the IBM **enterprise rack** offerings and other rack options

For More Information

IBM System x and xSeries Servers	**ibm.com**/systems/x
Electronic Service Agent	**ibm.com**/support/electronic
IBM System x and BladeCenter Power Configurator	**ibm.com**/systems/bladecenter/powerconfig
Standalone Solutions Configuration Tool	**ibm.com**/servers/eserver/xseries/library/configtools.html
Configuration and Options Guide	**ibm.com**/servers/eserver/xseries/cog
ServerProven Program	**ibm.com**/servers/eserver/serverproven/compat/us
Technical Support	**ibm.com**/server/support
Other Technical Support Resources	**ibm.com**/servers/eserver/techsupport.html

Legal Information

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

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

IBM, the IBM logo, the e-business logo, BladeCenter, Calibrated Vectored Cooling, PowerExecutive, System x, Virtualization Engine and xSeries 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.

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

Microsoft, Windows, Windows NT, Windows XP and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

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

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

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

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

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

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

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

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

XSW03010-USEN-00