EDISON Group
89 Fifth Avenue, 7th Floor
New York, NY 10003
www.TheEdison.com
212.367.7400

Blade Server Power Study

IBM BladeCenter and HP BladeSystem

November 7, 2007

First Publication: November 2007
Produced by: Craig Norris, Author; Luis Reyes, Senior Analyst; Andrew Podosenin, Senior Analyst; Barry Cohen, Editor-in-Chief

Table of Contents

Executive Summary

This report presents the results of tests conducted by Edison Group to compare the power consumption of the IBM BladeCenter blade server system with a comparable BladeSystem blade server configuration from HP. It provides background, configuration details, and methodology by which Edison reached its conclusion: that IBM BladeCenter H requires nearly 10 percent less power than the equivalently configured HP BladeSystem c7000.

Objective

Edison sought to address the lack of any industry study reflecting high-volume blade server configurations representative of those commonly deployed today within customer installations. We developed tests to establish parity between IBM BladeCenter and HP BladeSystem setups in a lab setting and measure the actual power consumption of each.

Many of the components within various commercially available servers —processors, memory, fans, power supplies, etc. — are supplied by the same manufacturers. For this reason, no appreciable difference in power consumption exists between servers at the component level. However, IBM BladeCenter as a whole (blades and chassis combined) is purposefully designed to maximize the efficiency of power delivery and cooling. Edison sought to accurately measure the power consumption of the IBM BladeCenter platform as compared to that of HP's BladeSystem. Therefore, we benchmarked a workload on complete assembled systems.

NOTE: IBM refers to the structure in a blade server that contains the blades as a "chassis," while HP refers to the same thing as an "enclosure." For the sake of clarity we use the term "chassis" throughout this report.

Methodology

The basis of this study consists of benchmark tests conducted using equipment configured within a test lab environment reflecting common data center practices today. Our tests were designed to assess power consumption of IBM BladeCenter solutions as compared to that of HP BladeSystem solutions. Two different comparison tests were conducted:

1. **Comparison One — Local storage:** The goal of this comparison is to determine power consumption using blade configurations deployed with local storage similar to rack servers.
2. **Comparison Two — Boot from SAN:** Booting from a Storage Area Network (SAN) is widely regarded as an industry best practice. Implementing diskless blade servers can provide increased administration efficiency, more flexible workload relocation, improved storage utilization, reduced disk costs through consolidation, and new opportunities for disaster recovery. The goal of this comparison is to also remove local disk drives from the server; this removes differences in power consumption by the local disk drives as a potential affecting factor in Comparison One.

For full details on the configurations used in testing, refer to the section entitled "*Configurations Tested*" later in this report.

Details of the test setup and procedures used in the comparison testing are provided in the section of this report entitled "*Study Methodology*."

Audience

This study will be useful to data center administrators as well as to any executive or manager overseeing an IT department as a cost center. It will be of particular value to those considering an investment in IT infrastructure and who must plan for the resources necessary to support it, either for simple reasons of practicality or to financially justify decisions in making that investment.

Summary of Findings

As can be seen from the results charts presented below, in both configurations, the IBM system power consumption per blade is considerably better in terms of electrical usage for a data center of enterprise scale. **In fact, IBM BladeCenter H requires nearly 10 percent less power than the equivalently configured HP BladeSystem c7000**.

System	IBM BladeCenter H Local Storage	HP BladeSystem c7000 Local Storage	IBM BladeCenter H Boot from SAN	HP BladeSystem c7000 Boot from SAN
Server Blades Per Chassis	14	16	14	16
Peak Power Consumption per Server Blade (Watts)	300.63	333.42	299.82	325.41
IBM Advantage - Percent Less Power Consumption Per Blade	9.84%	N/A	7.86%	N/A
Server Blade BTU/Hr	1025.14	1136.97	1022.38	1109.64
IBM Advantage - Percent Less BTU/Hr Per Server Blade	9.84%	N/A	7.863%	N/A
Combined Server and Cooling Power Consumption (kWh)	134.68	149.37	134.32	145.78
Combined (Server & Cooling) Cost per year Uniform Configuration (224 blades)	$110,902.12	$122,999.76	$110,603.41	$120,042.91
Combined (Server & Cooling) Savings per year, Uniform Configuration (224 blades)	$12,097.64	N/A	$9,439.50	N/A

NOTE: Detailed explanation and calculation methodology can be found in the "Test Results" section of the report.

An approximate 10 percent savings in power can be highly significant for many customers, especially for a large data center. To fully grasp this significance, extrapolate the number of blades to 224; this is the number required to have full chassis with an equal number of server blades for both IBM BladeCenter and HP BladeSystem. The official energy use statistics from the U.S. government [1] states that the average cost per kilowatt hour of electric power in the United States is 9.4 cents.

Using this figure with the extrapolated uniform configuration number of 224 servers, a 10 percent savings in energy costs realized by using IBM BladeCenter over HP BladeSystem systems can amount to approximately $12,000 per year.

[1] "*Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State*," published by the Energy Efficiency Administration of the Department of Energy.

Introduction

The "Green Revolution" has had builders of commercial facilities vying for certification as being environmentally friendly. However, it is not driven solely by altruistic concerns about global warming and dwindling sources of energy. Going green can have real and positive impacts on the cost of businesses making use of those facilities. This is especially true for IT data centers. Many high-level corporate decision makers are avidly seeking solid information about ways to adjust their data center planning and what to look for in the IT server market.

IT departments have lagged behind facility planning in adopting green technologies. However, trends in the data center have now begun to force the matter to a head. As the amount of computational capacity per physical space within data centers continues to climb, power and cooling requirements have emerged into the forefront of challenges facing data center architects. It is no exaggeration to characterize the situation confronting many organizations as a data center energy crisis. Many data centers have reached full capacity, limiting their organization's ability to grow and make necessary capital investments. They want the latest generation of servers but cannot use them because there is not enough electrical power and/or not enough cooling. In many urban centers, where space is at a premium and public utilities are strained to the limit, data centers may be severely restricted in their options regarding power and HVAC capacity expansion. There are locales where the municipal utilities themselves cannot accommodate requests for further electrical power to a given address, because the distribution grid or the generating capacity has reached its limit.

In all areas, whether urban or rural, rising energy consumption has also claimed the attention of those concerned with the bottom line. As the cost of many IT resources have leveled off or even fallen, the increasing cost of energy has come to represent a growing percentage of a data center's total cost of ownership. According to analyst firm IDC, for every dollar spent on computer hardware roughly 50 cents is spent on energy for power and cooling — an amount expected to increase by 54 percent to 71 cents per dollar over the next four years.[2] If trends continue on their present track, in the near future there will come a point where energy costs more than the servers themselves.

For all these reasons, the use of electrical power has become a critical factor in planning data center infrastructure. Increasingly, the power consumption rating for any given component of an IT infrastructure environment is only one of many interrelated facets to

[2] IDC, *Worldwide Server Power and Cooling Expense 2006-2010 Forecast*, Doc #203598, September 2006.

consider. Blade servers are positioned to help address the issues of power consumption and, accordingly, cooling requirements.

The goal of infrastructure design and planning is to take a more holistic approach in order to reduce the power consumption/heat emitted wherever possible. When applied to data center servers, this can encompass everything from planning the optimum location in terms of room airflow and HVAC resources, down to the design of a single blade server itself. For example, IBM internal studies show that an average IBM BladeCenter chassis with embedded Ethernet and fibre channel switches can help save customers up to 50 percent on power per port over a typical rack optimized server.

The Impact of Blade Server Systems on Data Center Power Usage

By allowing multiple individual and self-contained computer servers to share common infrastructure resources — such as power, cooling, networking, and management — of a single chassis, blade servers have enabled densities of 80 computers and more per industry standard 42U rack. Considering that the traditional standard server rack configuration is limited to a maximum of 42 computers, conserved use of space is an obvious benefit, and with real significance to the bottom line. It has been estimated that a data center space costing $20 million a few years ago can cost as much as $300 million today.[3]

Considerable power conservation can be realized through the pooling and sharing of common infrastructure in support of an entire chassis that can each now hold multiple discrete servers rather than a single one. The shared power and cooling of a blade server helps to reduce the amount of heat generated as compared to traditional rack and tower servers, and newer chassis designs feature sophisticated controls and adjustable cooling systems.

IBM has been a pioneer in leveraging the advantages in the cooling efficiency that a blade configuration allows; for example, BladeCenter was designed from its inception to utilize shared cooling at a time when most blade server vendors still employed the much less efficient approach of having a fan on every blade. Shared power, too, was an innovation that IBM introduced from the BladeCenter's inception — an approach that has since been adopted by other blade vendors. (The later section of this report entitled "*IBM's Approach to Power Efficiency*" provides details on the various engineering strategies IBM uses in optimizing blade server power and cooling efficiency.)

---

[3] "*The Green Data Center*," Matt Stansberry
http://searchdatacenter.techtarget.com/general/0,295582,sid80_gci1273283,00.html

With a dramatic increase in the concentration of computing power comes a dramatic increase in the use of electrical power within a given physical space. Some studies estimate that power requirements at the largest data centers have been growing at more than 20 percent annually. One such study estimates that the use of electricity associated with servers doubled between 2000 and 2005. [4] Real roadblocks loom in the foreseeable future as demand outstrips the capacity of traditional power generation to keep up, and as energy costs continue to climb.

Trends Contributing to Greater Concentration of Power Consumption

Certain trends contribute to an ongoing rise in demand for data center power as well. A steady increase in the speed and capacity of processors and a steady increase in the number of processors within each server are both examples of these trends.

Another example is the determination on the part of businesses to address the problem of underutilization of server resources. This problem received a great deal of attention in the industry media a few years ago. It accounts for the rapid adoption of virtual servers, using software products such as VMware, to consolidate still further. Consolidating workloads onto fewer platforms can help simplify and optimize existing IT infrastructures, including servers, databases, applications, networks, and systems management processes.

Virtualization also allows for greatly streamlined detailed tracking and management of hardware configurations, even in widely distributed systems. It provides centralized insight into and control over the usage and performance of critical components such as processors, disks, and memory. Increasingly, organizations are implementing virtualization solutions not only to further reduce application footprints and optimize server usage, but to also realize considerable savings in TCO by reducing administrative complexity and expediting high-availability strategies.

All these developments in the data center — while beneficial in themselves —contribute to the rising demand for more electrical power per physical space. It is something of a paradox that the more powerful server platforms become, the more work that is demanded of them. The result is that the demand for applications exceeds even the increased capacity of servers. This is one explanation for why the number of server units also continues to grow despite greater efficiency.

---

[4] "*Estimating Total power Consumption by Servers in the U.S. and the World*," Jonathan G. Koomey, Ph.D., Staff Scientist Lawrence Berkeley National Laboratory and Consulting Professor, Stanford University, sponsored by Micro Devices, February 2007.

Cooling to Processing: a One-to-One Ratio in Power Consumption

Apart from the direct consumption of electrical power to drive more computational processing within a given space, the additional draw on energy to power cooling and auxiliary equipment must be considered. It has been estimated that for every kWh of electricity used for processing IT loads, another is expended for supporting infrastructure components such as UPS, power distribution units, air handlers, pumps, chillers, and other devices. [5]

Comparatively, a given number of blade servers will consume far less power for both processing and cooling than the same number of standard 1U servers. However, because a rack of blade servers will nearly always contain far more servers within the same physical space, it will likely require more cooling capacity per rack.

Altogether, the increasing energy demands of the data center account for no small expenditure. Analyst firm IDC estimates that $29 billion was spent on power and cooling IT systems in 2006. [6]

IBM's Approach to Power Efficiency

IBM has formulated Cool Blue™ technology — a comprehensive portfolio of solutions combining efficiency, planning and control — to meet power efficiency demands of the data center using a holistic approach. Cool Blue provides smart, new product and data center designs and a method to control and monitor system power and heat requirements.

IBM takes the approach that successful handling of the complexities of power and cooling begins with an intelligently designed server that takes less power to deliver the function and performance needed. These servers are paired with a well-planned rack and data center layout, and used with power-management tools that measure power and heat while collecting and trending power consumption and temperature data server-by-server.

The IBM energy management portfolio tackles the challenge to increase power and cooling efficiency and reduce costs on four levels:

---

[5] *Ibid.*

[6] IDC, *Worldwide Server Power and Cooling Expense 2006-2010 Forecast*, Doc #203598, September 2006.

- **Inside the system —** All IBM BladeCenter servers start with Calibrated Vectored Cooling technology, which is a holistic design approach that provides efficient cooling while maintaining reliable operation of the servers. The methodology allows for optimization of system layouts and airflow while ensuring no wasted airflow for energy efficient cooling designs. Coupled with more energy-efficient power supplies, IBM BladeCenter servers can generate less heat in the critical AC-to-DC power conversion than many alternative systems on the market today.
- **In the rack —** IBM BladeCenter servers are designed to work at full density in a well–planned rack solution. They are also designed to operate at extended temperature ranges to keep the system up and running — even in extreme temperature and failure conditions. IBM rack solutions are engineered to optimize air flow and prevent undesirable air recirculation within the rack, meaning that the servers can run in optimal temperature conditions.
- **With IBM BladeCenter** —Consolidated infrastructure in the chassis results in energy savings that can be utilized for increased processing capability in the same power and coding envelope, as well as better utilization of floor space to "right size" the data-center design.
- **In the room with servers and storage** — The data center is key to a healthy server environment. If a data center room cannot handle the kind of density your business requires, smart rack-level heat solutions like the super-efficient IBM Rear Door Heat eXchanger should be considered.

The Intent of This Study

Edison sought to conduct the tests described in this study after noting the lack of any industry study reflecting high-volume blade server configurations representative of those commonly deployed today within customer installations.

Existing studies involved double-wide and triple-wide server configurations[7]. IBM does offer customers the flexibility to purchase expansion options; however, these options are typically used by customers in very unique cases and for specific applications. Ordinarily, customers deploy fully populated chassis with single-wide blade servers. It was felt that a study based on this more common scenario would be far more useful for end users and would appeal to a much wider audience.

[7] Both IBM and HP offer double wide and triple wide server blades. These servers typically contain additional processors, memory, storage and I/O and require additional chassis slots – thus the name.

Study Methodology

As described more fully in the Executive Summary subsection entitled "Methodology," Edison conducted two separate comparison tests: [8]

- **Local Storage Configuration** – The server blades in both chassis were configured with two 1.86GHz 80W Intel® Quad-Core (E5320) processors, 8 GB of memory in 1GB DIMMs, and internal storage.
- **Boot from SAN Configuration** – The same server configuration was used in this comparison, but the internal storage was removed and the systems were booted from a Fibre Channel SAN.

The subsections presented here describe specifics employed in the testing.

Power Consumption

To measure power utilization, Edison's lab metered [9] the two power lines plugged into the servers. The results were combined and averaged to account for very slight variations between them.

Test Design and Test Procedures

The test scenarios were designed to compare configurations from both HP and IBM. The components in these configurations were matched as closely as possible within the different system architectures. Wherever possible, identical components were selected — i.e., CPUs and memory. The objective was to attain relative power consumption between IBM BladeCenter and HP BladeSystem. Therefore, the processors were chosen with power consumption, rather than raw performance, as a priority. There are higher-performing processors available from both vendors, but the ones selected offered a good mixture of performance and power consumption. The other components in the chassis and blade servers were identical (or as close to it as possible).

[8] A word on processor choice: The 1.86 GHz processors were chosen because processors of that speed are available from Intel in both 50 Watt and 80 Watt designs. This speed was selected so that the SPEC JBB 2005 results would be as similar as possible for all of the server blades tested.
[9] The meter used is a Dranetz PowerXplorer PX5. Full meter specifications are in the appendices.

The one difference was the choice of internal storage for the local storage comparison. For that test, the new solid state drives were selected for the IBM BladeCenter tests while conventional disk drives were used in the HP BladeSystem[10]. Because HP offers no solid state drives, we selected the closest alternative.

Edison recognized that there would be a power utilization difference solely due to the drive choice; for this reason the "Boot from SAN" scenario was introduced. By using the same storage system for both server solutions, differences in storage configuration and internal storage power consumptions were eliminated.

Each server ran the following software:

Operating System	Windows 2003 Enterprise Edition with SP2 (32-bit)
Java Version	Jrockit-R27.3.1-JRE1.5.0_11
Performance Application	Specjbb2005 version 1.07

Device driver and other patches were applied as required for functionality. Configuration, provisioning, and deployment of the blade servers were performed using the respective vendor's management software and modules.

Once the servers were provisioned and tested, their clocks were synchronized and metered test runs were initiated using the SPEC JBB2005 run command line. The tests were timed to start concurrently through the use of the Microsoft Task Scheduler. Two test runs were performed, with no significant delay between runs.

SpecJBB2005 was chosen to generate a 100 percent processor load, and the test parameters were tuned to approach as near-equal performance from the systems as possible. Tests were not selected to measure maximum performance, but to achieve similar performance.

The power meter generates logs and reports of its results. The meter readings, made at one second intervals, were averaged for each channel. The results were then adjusted for peak power consumption. Power utilization during ramp-up and ramp-down was eliminated as the duration of these periods varied from test to test, affecting the averages.

[10] IBM is beginning to use solid state drives in its servers because the company's statistics show that 49 percent of server failures are due to hard disk problems.

The remaining results are based upon simple calculations:

- BTU per chassis was calculated by multiplying the Watts per Chassis by the standard BTU/Hr conversion factor of 3.41
- The comparisons between chassis were calculated by dividing the IBM results by the HP results.

Test Environment and Equipment Used

In addition to the computer and storage hardware listed elsewhere, the test environment for all systems included:

Temperature Controlled Test Room	• 75° F =±2 degrees
Server Racks	• APC Netshelter VX 25U chassis w/sides and panel blanks as needed
PDU	• Identical PDUs were used to connect six power cords from chassis to two circuits
Power Meter	• DRA/PX-5 - Power Analyzer • DRA/DRANVIEW 6-ALL Software • DRA/LEMFLEX3K;D 3000 AMP Flex CT
Miscellaneous	• NetGear GSM 7248 GB Ethernet switch • Keyboards, mice, and monitors as required

Default Chassis Hardware Tested

Default Chassis Configuration	IBM BladeCenter H Chassis	HP BladeSystem c7000 Chassis
AC Power Input	Single-Phase	Single-Phase
Full fan/blower config	2	8*
Full power supply config	4	6
Management	2 x Advanced MM	BLc7000 Management Module
Ethernet Switches	2 x Nortel Layer 2/3 Cu GbE	2x BLc GBE2c LY 2/3

* The 8-fan configuration for the HP chassis was recommended using the HP BladeSystem Power Sizer v2.7.

Server Blade Configurations Tested

Configuration Details	IBM	HP
Comparison One – Local Storage		
Chassis	IBM BladeCenter H Chassis	HP BladeSystem c7000 Chassis
Blade Server	14 x IBM BladeCenter HS21 XM	16 x HP ProLiant BL460 c G1
Processor	2 x 1.86 GHz 80 Watt Intel® Xeon® Quad-Core (E5320)	2 x 1.86 GHz 80 Watt Intel® Xeon® Quad-Core (E5320)
Memory	8 X 1GB DIMMs	8 X 1GB DIMMs
Memory Interleaving	Interleaved	Interleaved
Local Disk	Dual 16 GB solid state drives per blade	2 x 36GB 10K SAS HDD per blade
Comparison Two - Boot from SAN		
Chassis	IBM BladeCenter H Chassis	HP BladeSystem c7000 Chassis
Blade Server	14 x IBM BladeCenter HS21XM	16 x HP ProLiant BL460c G1
Processor	2 x 1.86GHz 80 Watt Intel Xeon Quad-Core (E5320)	2 x 1.86 GHz 80 Watt Intel Xeon Quad-Core (E5320)
Memory	8 x 1 GB DIMMs	8 X 1 GB DIMMs
Memory Interleaving	Interleaved	Interleaved
HBA	1 x QLogic Dual port 4 Gb FC HBA	1 x QLogic Dual port 4 Gb FC HBA
Local Disk	None	None
Disk Storage System	IBM System Storage DS3400	IBM System Storage DS3400
Disk Storage	14 x 73 GB SAS 10K HDD	14 x 73 GB SAS 10K HDD

Test Results

Edison's research shows that, while running equivalent workloads, the IBM BladeCenter H was considerably more power-efficient than the HP BladeSystem c7000 for both evaluated configurations. When averaged for per-blade power consumption, the IBM system was almost **10 percent better** with internal storage and almost **8 percent better** when booted from a SAN. This significant difference in per-blade power consumption under the same workload can extrapolate to a considerable power and dollar savings when projected across an entire data center of blade servers.

Notes on the Results

Consumption per Blade and Uniform Configuration

The HP BladeSystem c7000 and IBM BladeCenter H chassis contain space for — and were tested with — a differing number of blades. To properly compare the results and calculate real data center energy cost savings requires that the number of server blades for each platform be equalized. This equalization occurs at 224 blades — that is, 14 full HP BladeSystem c7000 chassis as compared to 16 full IBM BladeCenter H chassis. The per-chassis power and BTU/Hr results have also been adjusted to compensate for the differing number of server blade servers.

Cost per Year

The cost per year is based upon a simple calculation of the power consumed, in kilowatt hours (kWh), times the number of hours per day times the number of days in a year. Finally the result is multiplied by the nationwide average cost of electricity as reported by the United States Department of Energy (presently $ .094).

The calculation applied may be written out formulaically as follows:

(Watts * 24 * 365)/1000) * $ .094

The tables below summarize the results of the Edison testing.

Comparison One —Local Storage

System	IBM BladeCenter H	HP BladeSystem c7000
Server Blades Per Chassis	14	16
Peak Power Consumption for Chassis (Watts)	4,208.80	5,334.75
Peak Power Consumption per Server Blades (Watts)	300.63	333.42
IBM Advantage - Percent Less Power Consumption Per Blade	9.84%	N/A
Chassis BTU/Hr	14,352.51	19,191.51
Server Blade BTU/Hr	1,025.14	1,136.97
IBM Advantage - Percent Less BTU/Hr Per Server Blade	9.84%	N/A
Uniform Configuration (224 blades)		
Server Power Consumption (kWh)	67.34	74.69
Cost per year	$ 55,451.06	$ 61,499.88
Server Power Consumption Savings per year	$ 6,048.82	N/A
Cooling Power Consumption (kWh)	67.34	74.69
Server and Cooling Power Consumption (Watts)	134.68	149.37
Combined (Server & Cooling) Cost per year	$ 110,902.12	$ 122,999.76
Combined (Server & Cooling) Savings per year	$ 12,097.64	N/A

Comparison Two — Boot from SAN

System	IBM BladeCenter H Boot from SAN	HP BladeSystem c7000 Boot from SAN
Server Blades Per Chassis	14	16
Peak Power Consumption for Chassis (Watts)	4,197.46	5,206.51
Peak Power Consumption per Server Blades (Watts)	299.82	325.41
IBM Advantage - Percent Less Power Consumption Per Blade	7.863%	N/A
Chassis BTU/Hr	14,313.34	17,754.19
Server Blade BTU/Hr	1022.38	1,109.64
IBM Advantage - Percent Less BTU/Hr Per Server Blade	7.863%	N/A
Uniform Configuration (224 blades)		
Server Power Consumption (kWh)	67.15	72.89
Cost per year	$ 55,301.71	$ 60,021.46
Server Power Consumption Savings per year	$ 4,719.75	N/A
Cooling Power Consumption (kWh)	67.16	72.89
Server and Cooling Power Consumption (kWh)	134.32	145.78
Combined (Server & Cooling) Cost per year	$ 110,603.41	$ 120,042.91
Combined (Server & Cooling) Savings per year	$ 9,439.50	N/A

Additional Test for High-Efficiency Configuration

The tests described above, which Edison conducted to compare power consumption between IBM BladeCenter and HP BladeSystem servers, were specifically designed to establish parity between the two systems under testing. Having been impressed with the results on the part of the IBM BladeCenter H configuration, Edison decided to apply its resources to a final test. This one used an enhanced configuration that IBM recommends to deliver the greatest reduction in power usage for the most power-constrained customers.

For this test, the blades used were two 1.86GHz 50W Intel® Xeon® Quad-Core (L5320) processors and 2 GB DIMMS.

The charts presented here show the configuration used and the results obtained.

Server Blade Configurations Tested in High-Efficiency Test

Configuration Details	IBM
High-Efficiency Chassis with High-Efficiency Server	
Chassis	IBM BladeCenter E Chassis
Blade Server	14 x IBM Bladecenter HS21 XM LV
Processor	2 x 1.86GHz 50W Intel Xeon Quad-Core (L5320)
Memory	4 x 2 GB DIMMs
Memory Interleaving	Non-Interleaved
Local Disk	Dual 16 GB solid state drives per blade

Results of High-Efficiency Configuration Testing

When a data center is being configured for particularly power-constrained environments, IBM would recommend the highly power efficient BladeCenter E chassis. In addition, the customer would be advised to select low-voltage processors and fewer memory DIMMs. **All of these changes combine to deliver up to 25% more power efficiency than the IBM BladeCenter H configuration with local storage — which already demonstrated *10 percent* greater efficiency than HP BladeSystem c7000**. The following table shows Edison's test results from running the same workloads described earlier on the high-efficiency IBM BladeCenter E Chassis instead. The comparisons are intended to demonstrate the further advantage available within the IBM product line.

System	BladeCenter H Local Storage	BladeCenter E 50Watt Local Storage
Server Blades Per System	**14**	**14**
Peak Power Consumption for System (Watts)	4,208.80	3,147.83
Watts Per Blade	300.63	224.84
Per Cent Power Consumption less than BladeCenter H	**N/A**	25%
System BTU/HR	14,352.00	10,734.10
Server Blade BTU/Hr	1,025.14	766.72
Per Cent BTU/Hr Less than IBM BladeCenter H	**N/A**	25%
Uniform Configuration - 224 Server Blades (16 IBM BladeCenter Chassis)		
Server Only Power Consumption (kWh)	67.34	50.37
Server Power Consumption Cost per Year	$55,541.06	$41,472.77
Server Power Consumption Savings over IBM BladeCenter H (kWh)	**N/A**	**16.97**
Cooling Power Consumption (kWh)	67.34	50.37
Combined Server and Cooling Power Consumption (kWh)	134.68	100.73
Combined Server and Cooling Power Consumption Cost per year	$110,902.12	$82,945.54
Combined Server and Cooling Power Consumption Uniform Configuration Cost per Year Savings over IBM BladeCenter H	**N/A**	**$27,956.58**

Conclusions and Observations

As can be seen from the results comparing 224 server blades per platform presented below, the IBM system utilized less electricity both for powering the system and for cooling the data center. In an enterprise-scale data center — where servers number in the hundreds or thousands and where the ratio of performance to power for infrastructure, cooling, etc. is approximately 1:1 — that difference in the cost of electricity is quite significant.

System	IBM BladeCenter H Local Storage	HP BladeSystem c7000 Local Storage	IBM BladeCenter H Boot from SAN	HP BladeSystem c7000 Boot from SAN
Combined Power Consumption (kWh)	134.68	149.37	134.31	145.78
Combined Cost per year	$110,902.12	$122,999.76	$110,603.41	$120,042.91
Combined Savings per year	$12,097.64	N/A	$9,439.50	N/A

Furthermore, where saving on power costs is a particularly compelling goal, the power-optimized IBM BladeCenter E is recommended. When configured with lower power-consuming processors such as the 1.86GHz 50W Intel QC (L5320) and fewer memory DIMMs, it can yield savings that are even more significant. Edison's tests demonstrate that such power-optimized blade implementations can deliver the same performance for many applications as their greater power-consuming counterparts. This makes them an ideal choice when expandability could otherwise be achieved only with the more power hungry systems.

System	IBM BladeCenter H Local Storage	IBM BladeCenter E
Cooling Power Consumption (kWh)	67.34	50.37
Combined Server and Cooling Power Consumption (kWh)	134.68	100.73
Combined Server and Cooling Power Consumption Cost per year	$110,902.64	$89,945.54
Combined Server and Cooling Power Consumption Cost per Year Savings over IBM BladeCenter H	N/A	$27,956.58

The savings of almost $28,000 per year over the IBM BladeCenter H and even more over the HP BladeSystem C can have a profound effect of data center operating costs. Data center planners who select high-efficiency server and other intelligently designed data center infrastructure technologies can save their organizations significant operating costs while lowering their carbon emissions footprint.

Observations

This final section presents various issues encountered and observations made by Edison analysts while setting up and running the test equipment. It should be noted that both companies offered excellent phone support.

IBM BladeCenter

One thing that our analysts noted is that IBM offers a better virtual KVM solution with a built-in KVM port for direct connection. It allows for plugging in a monitor/keyboard/mouse directly to the system to access all 14 blades through a simple switching between the screens. HP's system requires manual plugging in of a cable to each blade to be accessed directly. While this is not relevant for normal operations over a network, direct connection is desirable for setup and configuration. The IBM systems also ran noticeably more quietly than did the HP systems.

HP BladeSystem

While the quality and professionalism of HP's onsite support was superb, we spent two weeks grappling with technical issues involving the HP setups. The notable issues are enumerated here:

- The embedded SCSI Array (E200) could not initially be detected on several blades. The resolution was to re-seat each card to the motherboard, with resolution taking one day. HP has a customer advisory posted on this issue.
- The Brocade fibre switch lacked the default username and password assignment. To resolve the issue, we needed to reset the switch. Resolution required one day.
- Six blades had defective HBAs. The World Wide Name on the card did not match the QLogic bios. Changing the HBA cards resolved the issue.
- No blade on Slot 10 of the chassis would boot from the SAN. The resolution required replacement of the back plane of the chassis.
- The HP systems also ran noticeably louder than the IBM systems.

The takeaway from Edison's experience with the test setup — particularly illustrated by our experience with the HP system — is to thoroughly test all functions that may be feasibly considered, even for *future* use. This lesson is drawn from the second to last bullet point mentioned, where the entire chassis backplane needed to be replaced.

Up until the "Boot from SAN" test was done — at a later date than all the other tests — it was felt that all issues had been addressed and that things would run without a hitch. Our experience is a forceful reminder that even serious problems can remain concealed unless revealed by thorough testing upfront. Even with a serious and potentially disruptive problem, a blade server system in a production data center could well be up and running for a long time until a reconfiguration brings it to light. This could potentially occur after the manufacturer's warranty has long since expired.

It is also important to note that the problems we experienced with our HP BladeSystem should be recognized as unique. We do not believe that HP is in the practice of shipping defective equipment, nor do we believe that our channel source regularly ships hardware with problems.

BLL03002-USEN-00