



ROBERT FRANCES GROUP
Business Advisors to IT Executives

120 Post Road West, Suite 201, Westport, CT 06880 Phone: 203-429-8950 Fax: 203-429-8930

Increasing Energy Efficiency with x86 Servers

Jerald Murphy
Senior Vice President and Service Director
Robert Frances Group



Executive Summary

The current economic crisis has most companies scrambling for ways to get the most out of every dollar they spend. This process leads to increasing uses of automation, which adds costs to IT at the same time business are looking to rein-in costs. Many IT executives have moved to consolidating data centers and servers with virtualization and blade servers. However, to date, many of the gains have been illusory due to high server costs, lack of effective virtualization and management, and inefficiencies that make it expensive to power and cool these consolidated systems.

Most large enterprises are now very familiar with the problems associated with escalating power and cooling costs in the data center. As more business processes have become automated, storage requirements have increased as well. This additional centralized capacity has resulted in a significant rise in energy and cooling costs for these systems.

Over the last five years, the total power generated by individual compute platforms has grown by over 100 watts, with system power consumption increases of over 300 watts per server. The energy needed to power and cool data centers cannot be ignored by IT, since the total energy needed over the life of the server often exceeds the cost of the systems purchased. The current economic situation further exacerbates measures, since more money is not available to improve cooling, and power densities often cannot be maintained. RFG has seen that power efficiency has become the number one IT issue in the data center, and will continue to be an important factor for the foreseeable future. In addition, the ability to manage these new virtualized environments will take on an increasingly important role. IT executives need to anticipate and adapt future data center designs and operations to these new challenges

RFG believes that IT executives are going to get the best chance of optimizing power, cost, and performance by selecting computing solutions that are designed to synergistically enhance capabilities in these three areas. IBM Corporation has teamed with Intel Corporation to integrate the benefits of energy efficient computing. On the one hand, Intel has added capabilities in its Intel® Xeon® Processor 5500 Series architecture that increase computational throughput for the same or less power. In addition, its integrated memory management has improved data throughput for less power as well. IBM has extended the capabilities of the Intel Xeon Processor 5500 Series based System x and BladeCenter offerings. They have improved memory management, adaptive energy management, and better virtualization. This helps get the most processing out of each data center element for the least cost. For example, HS22 offers outstanding performance, flexible configuration options and simple management in an efficient server designed to run a broad range of workloads. The new IBM BladeCenter HS22 offers 60 percent better price/performance¹ on virtualization and 33 percent better on general business

¹ 60% price performance is based on SPECint_rate_base projected performance improvements comparing the HS21 using top-bin 3.33GHz 5400-series processors compared to the HS22 top bin 5500-series 2.93GHz processors. Pricing comparison based on standard list price with HS21 using 1 processor and 1GB memory versus HS22 standard list price for mode 7870-C4x, which has 1 processor and 2GB memory



applications. It also offers 50 percent more memory capacity² and 12 percent less power³ used by the processor when fully utilized (50 percent less when idle) as compared to its predecessors. In addition, the HS22 is backwards compatible with all BladeCenter enterprise and office chassis, including the original BladeCenter E so those who have a current chassis installed can fill empty slots without having to rip and replace hardware or sacrifice performance. IBM also offers powerful management tools such as IBM Systems Director Active Energy Manager™, which allows customers to more fully understand and manage their energy usage by monitoring, trending and capping of actual server power consumption. The new features of the new generation servers, combined with IBM's proven management and automation capabilities, deliver the maximum performance per watt and dollar in the data center.

While Dell and HP will certainly implement solutions with the Intel Xeon Processor 5500 Series architecture as well, their less efficient chassis (e.g. less efficient power supplies, more air handlers) and lack of UEFI for more complete remote management will not offer the same compelling value as the IBM/Intel combination. With Cisco Systems, Inc.'s nascent entry into the blade server market, RFG also expects that it will take Cisco considerable time to establish the necessary management capabilities to deliver the value IT executives are looking for in next-generation x86-based blade server architectures.

Overview of Current Data Center and Business Realities

Over the last two years, IT executives have been confronted with a significant shift in their operations costs. These shifts have been most noticeable in the data center, where more compute power and increased server density have led to a dramatic increase in power consumption.

Energy Issues in Current Data Centers

In an effort to get a better handle on distributed IT resources and perceived inefficiencies, many IT shops have moved to reduce the number of data centers in order to increase the efficiency of compute resources. Stand alone computers typically run at 15 percent CPU utilization and 20-30 percent storage utilization. Centralizing these resources has allowed IT managers to share compute resources to increase the utilization of servers. At the same time, the continued decrease in the size of integrated circuits and increased clocks speeds has increased server performance. While the trends of data center consolidation and increased compute density have increased the transactions per dollar spent on computing, these trends have brought with them their own problems with space, power, cooling, and management.

When there is space for these servers, the power characteristics of many of the current servers is not conducive to efficient operations. Many of the current series of blade servers have power densities that prevent a full rack of blade servers from being

² HS21 XM has 8 DIMMs. HS22 has 12 DIMMs. Each blade supports max 8GB memory DIMMs. 12 = 50% more than 8.

³ Intel internal measurements using SPECjbb2005* as of Aug 2007. System configurations: 2S, 80W processors, 8 DIMMs, 1 HDD, 1 PSU. Except Irwindale=110W processors; Power measurements using SPECjbb2005. Tylersburg-EP from Intel internal measurements as of Sept 2008 with 2.93GHz 95W processors This information is preliminary and subject to change before launch.



populated. Most data centers can handle power densities up to 15 KW, where some blade servers would populate racks requiring up to 30KW of power. Not only do most power distribution units (PDUs) not distribute this level of power to a rack, most existing data centers cannot dissipate the heat generated by servers with this energy profile. The net effect here is that available space in racks cannot be used, making available data center space much less in most real world scenarios.

In addition, there is a lack of effective power management in most of these servers. With typical server utilization rates of 10 percent, higher densities could be achieved if the power provided to servers could be modified so less power and cooling were required at lower workloads. In fact, power needs to be used more efficiently across the entire spectrum of workload.

Server Problems

The servers that exist in most data centers today have been designed for performance and cost optimization, not energy efficiency. With some companies finding up to 40 percent of their operational costs going to providing power and cooling to these systems, the cost of energy over the server's life can be more than the purchase price of the servers. Inefficiencies fall into several areas, including power conversion, utilization, and systems management.

Many servers in data centers today have power supplies that are only 70 percent efficient. This means that 30 percent of the power going to the server is simply lost as heat. Having inefficient power supplies means that excess money is being spent on power with additional cooling needed as well. Another problem with current servers is that they are used at only 15 percent of their capacity. A big problem with this from a power and cooling perspective is that the amount of power required to run traditional servers does not vary linearly with the utilization of the server. That is, ten servers running each at ten percent utilization will consume much more power than one or two servers that each run at 80-90 percent utilization.

In addition to these inefficiencies are the power densities of some of the early blade servers that came onto the market. Blade servers can consume up to 30 kW per rack, which either cannot be cooled, or requires special cooling systems to achieve these densities. The cooling systems that can achieve this density need to be installed in adjacent racks, which take up precious data center space. Some systems can augment cooling by going above the rack, but these can only work if the room has sufficient height to allow the installation of these systems. Some server manufacturers even have rear door heat exchangers. However, RFG cautions that some of these systems require customized racks to work, and may also not fit if the depth required is beyond what the existing rack topology will allow. IT executives should make sure any supplemental cooling technology can fit within the envelope of the current data center architecture, or make sure that future data center plans allow for the space needed by these technologies.



Finally, a lack of effective power management is often a problem with many existing systems. It would be great if the power could be scaled back if the systems did not need the power. However, many of the existing systems do not have the ability to change the voltage or frequency of the processors, and often cannot even give the IT manager insight into actual system power consumption. Very often, data center managers must rely on facilities managers to get an accurate handle on actual power consumed by data center systems.

Economic Issues in Business

The current economic crisis is exacerbating these data center issues. Most IT managers are finding that they have little or no money available for investment, without the ability to show hard return on investment (ROI). This is causing many IT managers to get creative in financing projects, often moving to self-funded projects. This money squeeze is occurring in all areas, both in capital dollars as well as operational costs.

The cost of energy continues to rise, with the only current respite being the cost of real estate. RFG is seeing many companies looking to relocate data centers to geographic regions that have lower real estate costs, lower energy costs and favorable labor rates. The current pressure to lower personnel costs has put tremendous downward pressure on staffing, requiring systems that are easier to implement and manage.

Actions that Need to be Taken to Fix These Problems

RFG believes there are many actions that can be taken to address the current and upcoming data centers. Actions should be taken by both data center managers as well as the systems vendors.

Data Center Manager Actions

Data center managers can take actions today that will mitigate most of these data center cost and environmental pressures. The best actions data center managers can take to address these issues are to implement intelligent server consolidation, with methodical virtualization of compute resources, including proactive planning between facilities and data center management. Intelligent refresh of data center servers can help bring about gains in most of these areas. Personnel costs can be addressed by increasing the automation and control of data center servers.

Server Consolidation

Server consolidation is the first step in improving IT resource utilization. With typical servers often at ten percent utilization, it makes sense to have a single server handling more than one process. This consolidation has to be analyzed, planned and executed carefully, in order achieve the optimal results for all IT and other business units involved.

Where individual business units have purchased separate servers, some decisions may need to be made on who "owns" and "controls" consolidated server resources. For this reason, it is best to consolidate similar applications being run by the same business unit, where possible. Attention should also be played to databases running on servers, and look



to consolidate similar databases onto consolidated servers. Another often-overlooked issue on consolidation is to reduce the servers needed by eliminating unneeded applications and databases. Server sprawl is often the result of a "one application/database - one server" philosophy that leads to this proliferation. While virtualization and consolidation can solve the immediate physical sprawl, effective systems planning and management are needed to avoid falling back into this trap. This can be done through application portfolio management. Many large companies run hundreds, even thousands of applications. Often times, legacy applications are running that are no longer needed because of mergers, acquisitions, and the change in business processes. IT executives who consolidate servers intelligently can dramatically reduce current and future data center power requirements.

Virtualization

RFG has seen a lot of hype around virtualization, and virtualization can mean different things to different people. Put simply, virtualization allows logical separation of an IT service from the physical system that provides that service. Many people think of virtualization today in terms of servers, but networks, storage, and practically any IT service can be virtualized.

For servers, adding a virtualization layer is what allows putting multiple applications running on different operating systems (as well as on different servers) to exist on a single piece of hardware. RFG suggests IT executives look closely at virtualization as a way to consolidate data center assets. Most IT shops start virtualization projects on non-mission critical systems, gradually working more applications onto a virtualized platform over time. IT managers should recognize that adding a virtualization layer to a system to allow consolidation does add complexity, which can add to management and security challenges. Therefore, IT executives should look for system solutions that are designed to simply adding virtualization to the IT infrastructure.

Increased Automation and Control

IT executives need to increase the automation and control of servers systems. As systems become virtualized and increasingly dense, there will be less ability to directly control servers in the data center. In fact, for security purposes, data center managers want to reduce the number of people directly accessing systems in the data center. As applications and systems become more dynamic, the automation of provisioning new servers and dynamically migrating images from one server to another will become essential for effective operations. Using servers that allow standard interfaces into these systems will enable better scaling of automation, which will help increase the engineer to server ratio, reducing personnel costs while increasing data center flexibility and security.

Vendor Actions

While data center managers can take some actions now to reduce the power and space problems, more actions need to be taken by vendors to shift from a performance optimization focus to maximizing the efficient use of energy, cooling, and space.



More Efficient Servers

For energy efficiency, vendors should focus on making the processors consume less energy, as well as using the energy delivered as efficiently as possible. This is best done when there is close coordination in system design between the processor and server manufacturers.

RFG has seen many systems vendors start to focus on making better use of the silicon that is available, with a shift from high clock rates and voltages, to adapting the clock rate and voltage based on the workload presented. Another action vendors are taking is to increase the overall work throughput by increasing the number of cores on the processor, while lowering the clock speed and voltage. Processor vendors need to make intelligent architecture choices in order to deliver the most energy efficient performance. In this way, each processor runs more efficiently, while increasing the overall work performed per processor.

Not only does the basic power efficiency need to improve, but also vendors need to increase the adaptability of the server to adapt to different workloads. That is, servers need to be designed to use only the minimum amount of power necessary to do the work at hand. Memory management needs to be improved as well. Increased use of memory will reduce the latency of systems, as well as increasing the number of virtual images that can be run on a single physical server.

Better Virtualization Technology

As companies move to consolidate infrastructure into larger data centers, physical space needs to be optimized. One of the ways to do this is to put multiple individual applications, previously running on separate servers, onto a central blade server infrastructure. Unfortunately, many of these applications are often running on different servers with different operating systems. In order to allow these applications to exist on one platform, virtualization technologies are used to enable this consolidation.

While there has been a great deal of growth in virtualization capability to date, much of the existing virtualization technology is relatively static. Virtualization technology needs to be used that allows more autonomic, automated capabilities to deploy workloads and reallocate server assets as needed based on changing business conditions.

Better Power Management

Finally, vendors need to continue improving their focus on managing the power used by servers. As energy costs move to eat up as much as 40 percent of the IT operations budget, finding every area where efficiencies can be increased is essential. In fact, efficient use of power may be the single biggest factor in determining server TCO. Some of the areas that IT executives should look for vendors improving power management include power distribution efficiency, CPU power consumption, cooling efficiency, and adaptive power management. As companies move towards using blades for virtualization, having the ability to change the power consumed by a server based on its workload will become increasingly important.



The IBM / Intel Xeon Processor 5500 Series Combination Addresses these Issues

Fortunately, for IT executives, IBM and Intel have been working in close coordination to address many of these data center issues. Intel has made some significant advances with its new Intel Xeon Processor 5500 Series architecture that maximize both power and performance. At the same time, IBM has increased its System x efficiencies, incorporating Intel Xeon Processor 5500 Series architecture advances with its own improved energy-efficient architecture. The IBM BladeCenter® exemplifies this. For example, IBM BladeCenter, featuring Intel Xeon processor based technology, uses up to 50 percent⁴ less floor space, 35 percent less energy⁵, has up to 65 percent less in connectivity costs⁶, and uses up to 84 percent less cables than competitor rack servers without sacrificing performance⁷. The combination of IBM BladeCenter and Intel works together to optimize overall data center performance, which lowers operational and energy costs.

Intel Xeon Processor 5500 Series Architecture Efficiencies

Intel has made significant strides in addressing energy issues with its latest Xeon processor 5500 Series. The 45nm Intel Xeon processor 5500 Series based processor combines new power management capabilities, along with significant processor performance improvements, especially in the area of memory management.

Different Power Options

First, there are different processor versions from which power options that can be selected, depending on the desired server performance characteristics. There are Intel Xeon Processor 5500 Series that operate using very low power, as low as 60 watts, where the focus is on total reduction of power costs. Where high performance is required, the highest performing processors can deliver over twice the CPU performance at the same power envelope as previous generation processors, with a maximum power consumption of 95 watts.

P States, T States

The new Intel Xeon processor 5500 Series based processors can operate in different P States and T States, which give server vendors the capability of dynamically changing the performance and energy use of processors, so that the energy consumed by individual processors can be dramatically lowered when there is no workload. For servers at low

⁴ <http://www-03.ibm.com/systems/migratetoibm/systems/bladecenter/> IBM BladeCenter provides up to 84 servers with BladeCenter E vs 42 1U servers in a single rack and eliminates the space requirements for network switches.

⁵ *ibid.* Based on IBM power engineering test data. Blade power is average power of total chassis solution.

⁶ Based on IBM vs. HP web pricing

⁷ Based on calculations of 1 rack of (6) fully populated BladeCenter E chassis vs. 1 fully populated rack (42) of competitive 1U servers



utilization rates, this will offer tremendous energy savings. The increase of energy states from 3 in the Xeon 5300 series to 15 in the Xeon 5500 series (Nehalem), allows a drop in idle power from 50W to 10W per core, will also reducing the power state transition time from 10 microseconds to less than two microseconds. These capabilities will allow end users to see a dramatic decrease in overall system power consumption without adversely affecting system performance.

These capabilities can be tied to external energy management systems, so server vendors that properly tie into these interfaces can continuously evaluate the server energy characteristics, either capping the total energy used by the processor, or changing the energy characteristics as the server demands change.

On-Chip Front-Side Bus

One of the biggest advances for the Intel Xeon Processor 5500 Series based processor is placing the front-side bus on processor. This advance will dramatically improve the performance of system memory, lowering memory I/O latencies, and increasing the total addressable memory that can be placed on systems. All servers using Intel Xeon Processor 5500 Series based processors will use DDR3 memory, with the new architecture capable of handling more memory slots. The increased memory bandwidth has increased up to 3.63 times previous processors⁸, which will be especially valuable for applications that require high memory throughput, such as database applications⁹.

Significantly higher performance for same power envelope

The new Intel Xeon processor 5500 Series based processors will increase performance of previous processors up to 2.4 times, with two threads available per core for eight threads per processor. In addition, the new processors will add Intel's Turbo Boost technology, which can allow "over clocking" of individual cores, while keeping the total power used by the system within established environmental parameters. In addition, Intel has incorporated its Intel® Intelligent Power Technology, which minimizes energy use when maximum performance is not needed.

In addition, the new Intel Xeon processor 5500 Series based processors have much better idle system performance, which have reduced idle power consumption by 47 percent compared with its previous Xeon 5400 series processors. New Integrated Power Gates and Automated Low Power States allow individual cores to drop to almost zero power consumption, independent of work that is being done by other cores. This control of core power can be done either manually or automatically, dramatically increasing system power flexibility.

The new Intel Xeon 5500 Series processors also improve virtualization capabilities, with hardware assist that will improve the ability to use virtualization hypervisors and improving live migration of virtual operating system images and applications. Intel has

⁸ Results with Nehalem-EP processor at 2.93Ghz compared with Xeon X5482 at 3.20Ghz

⁹ 2.54x improvement in TPC-C Oracle benchmark, Nehalem-EP v. Xeon X5460 at 3.16Ghz



included VT-x on the processor, VT-d for directed I/O on the processor, with VT-c for improved network connectivity. Thus, Intel gives IT managers a holistic approach to virtualization that can improve virtualization efficiency across the entire system.

IBM System x Efficiencies

The new System x servers are designed to expand on the capabilities of the Intel Xeon 5500 Series processors to make the most of energy efficiency, memory, space, and management.

Design

First, the System x is built fundamentally as an energy efficient architecture. IBM is the only server manufacturer that is using the latest UEFI 2.1 standard. This is the next-generation basic input/output system (BIOS), which increases the efficiency and flexibility of the entire server system. The physical design of the servers is optimized to require fewer fans while still providing all necessary cooling. The real goal in fan management is to reduce the airflow to the least necessary to cool the processor and all other system components. This is done with variable-speed fans, which can dynamically change the spinning rate to deliver the precise airflow needed in response to changing temperatures.

The x3550 M2 and x3650 M2 servers also incorporate an altimeter, which allows fan speed to be adjusted, based on air density—if the system is running at sea level the fans can run slower than if at high elevation. This dynamic flexibility is important, since the energy utilization of a fan varies exponentially with the fan speed. In addition, the iDataPlex architecture has been designed with dramatically improved airflow, reducing the power required to extract server heat and lowering the cooling required in the data center.

Efficient power supplies are essential for saving money. The new IBM System x servers are optimized for AC/DC power, with power efficiencies up to 93 percent. While some systems in data centers have moved to large rack-based DC power rectifiers, IBM has found doing this does not save energy, since additional DC power conversions are needed throughout the rack to provide different DC power to the various system components.

Use of Memory

IBM has the best use of memory of any of the system vendors. By moving to 8 GB DDR3 DIMMS, the total memory available to the server has grown to 96 GB. The memory increases to 128GB including the System x rack servers. This will significantly increase the effectiveness of servers for virtual, as more memory will increase the number of virtual servers that can be run on a given system.

Advanced Power Monitoring and Management

IBM provides customers with software tools to help in the monitoring and management of power usage. IBM Systems Director Active Energy Manager™ is a systems management plug-in that monitors and trends the actual energy usage of IBM systems.



Accurate measurement of power usage is essential to becoming more efficient. Active Energy Manager can also help manage power by capping usage by server or group of servers and extends the scope of energy monitoring to include select facility providers to enable a more complete view of energy consumption within the datacenter.

Additional thought for design with an emphasis on power and cooling approach of some blade servers can lead to less points of failure as well. For example, IBM BladeCenter can use four power supplies to redundantly power up to 14 servers and associated switching. To do this in a typical 1U stacked solution would require well over 28 power supplies. IBM BladeCenter is also a great example of this forward thinking in a shared thermal solution- built with two hot swap blowers that consume only about 100W. These two parts replace the 112 fans typical in other servers that consume well over 600W. Thus, the similar cooling function can be performed with both increased reliability and dramatically reduced power.

The new IBM BladeCenter HS22 system fully incorporates the use of Intel Xeon Processor 5500 Series architecture turbo capabilities, so that the IBM BladeCenter HS22 system can turn turbo mode on or off, based on application and system demands. The IBM BladeCenter HS22 can also make use of processor over-clocking, without exceeding the maximum power specifications.

Space

The nine U height of the IBM BladeCenter H chassis allows IBM to put more systems into a rack, increasing the total performance capability per rack in the data center. The new IBM BladeCenter HS22 systems also are able to take advantage of any existing blade chassis investments, as the IBM BladeCenter HS22 blade will fit in both the BladeCenter H and BladeCenter E chassis. The 7U BladeCenter E chassis, allows 14 HS22 blades per chassis, for up to 84 servers per standard rack. When loaded with fourteen new, power-optimized BladeCenter HS22 dual processor Xeon 5500 series blades per chassis, IBM achieves a leadership density of 168 CPUs per standard rack (42U). In addition, the memory controller on the processor allows up to 50 percent more DDR3 memory to be put on a server. These new capabilities increase the flexibility that is given to IT managers to configure server/memory combinations to fit their specific business needs.

Space benefits are even greater with the iDataPlex. Its highly efficient liquid cooling actually eliminates the need for air conditioning, with exit temperatures that can actually be lower than front air temperatures. The iDataPlex rack level liquid cooling is 70 percent more efficiently than traditional air-cooling. This energy efficient eco-system can save up to \$10,148 per rack per year. In addition, the iDataPlex provides 50 percent more density to customers running out of data center space, within traditional infrastructure power and cooling limits. Also, with the Rear Door Heat eXchanger, the room can actually be cooled with the system.



Multiple Disk Options

IBM gives IT managers the ability to put server different disk options onto a server. The standard disks can be used, which consumer 10 watts, but cost the least. Systems can also be configured without disks, which increase the reliability of individual blades, when configured with the option of booting the server from the storage area network (SAN). RFG is seeing more IT managers use this as an option, to add flexibility to the environment, while increasing the mean time between failure (MTBF) of individual blade components.

Finally, there is an option to buy servers with solid-state disks (SSDs). SSDs consume only one watt vs. 10 watts, making them a much better option from an energy conservation standpoint. Additionally, the I/O capabilities of SSDs are 100 times better than spinning disks. While the cost for SSDs is currently 10x of spinning disks, the 100x I/O rate means that the cost per transaction for high-throughput applications will actually be less with the SSD server. Since the cost of SSDs is coming down at a much faster rate than the unit cost of spinning disks, RFG believes the combination of higher performance, lower cost per transaction, and higher long-term availability will make blade servers with SSDs a compelling combination for many data center applications.

Improved Virtualization Capabilities

IBM's System x servers support most of the leading hypervisors, including ESXi embedded models. The large available memory mentioned above, will help increase the number of virtual servers that can be implemented per physical server. This is significant, since each real server than can be virtualized directly saves 300-1000 watts.

In addition, the IBM servers are designed to take advantage of all the Intel Xeon Processor 5500 Series architecture virtualization capabilities, including VT-x, VT-d, and VT-c. Finally, the large I/O capacity of IBM servers is essential for handling the increased connectivity requirements that will be brought about by massive virtualization consolidation. Each virtual server will have its own requirements for external and cross-server connectivity. As such, the number of virtual servers that can be put on a single system will be limited by the amount of system I/O capacity. .

If IT managers need high-end performance that requires more than four cores, IBM has systems that take advantage of Intel's six-core 7400 series processors. The IBM System x3850 M2 and x3950 M2, further increase performance with a custom chip set designed by IBM that includes a "Snoop" filter to further improve the efficiency of system cache, enhancing scalability. IBM's Snoop filter enables the processors to determine cache entries, eliminating traffic on the front side bus. This allows processors to go directly to cache or main memory, dramatically increasing performance¹⁰.

¹⁰ The System x3950 M2 is one of the fastest high end Intel based servers in the industry, outperforming the competition in the x86 TPC-c benchmark by more than 50%. In fact, the IBM System x3950M2 is the first single x86 server to break the 1 million transactions per minute barrier on a TPC-C benchmark <http://www-03.ibm.com/press/us/en/pressrelease/25129.wss>

The way these high-end systems implement memory makes them the only 4U systems that can scale from two up to 16 sockets, for a total of 96 cores and 1TB of memory. The ABX processor allows these IBM systems to achieve the same memory speeds with higher line rates, lower wattage, at 1/4 the amount of watts than competitive systems. In addition, these systems use IBM's Xcelerated Memory Technology, which allows the use of more DIMMS at higher frequencies.

For virtualization to be effectively used by customers, a sound systems management strategy is required for both physical and virtual resources. Just as server sprawl can occur with physical servers, inadequate systems management can lead to virtual server sprawl. IBM has integrated virtualization management into its IBM Systems Director management product. Administrators can use the same tool to discover, inventory and monitor health of virtual and physical servers as well as their associated topology. Systems Director also provides linkages to VMware virtualization management to provide automation through event action plans which can trigger VMware actions such as Vmotion.

Availability Benefits

As companies move to consolidate more applications with virtualization onto blade servers to increase system utilization, many IT managers are becoming concerned about the availability of these systems. Any one blade system going down now has a more dramatic effect on the business. Not only is one blade failing going to effect more applications, but one chassis failing can impact a dozen or more servers. For these reasons, blades must be designed for high availability, or else IT managers have to pay additional money for redundant servers to achieve this availability.

In the case of IBM BladeCenter servers, the chassis have been designed with this high availability in mind; this capability is not available on the HP systems. In fact, HP has had a problem with its power supplies due to this design issue¹¹. IBM uses redundant power connections to allow independence from the energy source. This is augmented by a redundant power bus, also missing from the HP systems. IN RFG's opinion, the redundant power bus is especially critical, since failed power delivery to the blades can affect multiple blades. By employing a redundant power bus, uninterrupted power delivery is further assured, without having to invest in redundant servers.

Network connectivity is equally important to system redundancy, especially given the distributed and connected nature of today's business applications. To this end, IBM has a redundant mid-plane and redundant I/O, to ensure that connectivity is maintained to all servers and processes.

11

http://h20000.www2.hp.com/bizsupport/TechSupport/Document.jsp?objectID=c01519680&dimid=1012424238&dicid=alr_sep08&jumpid=em_alerts/us/sep08/all/xbu/emailsubid/mrm/mcc/loc/rbu_category_AdvisoryCUSTOMER/alert



Limitation of Other Platforms

While other platforms certainly will use Intel Xeon 5500 Series processors to improve processor efficiency, the existing architectures of these products do not have the same synergistic benefits that the IBM/Intel Xeon Processor 5500 Series combination offer.

HP Platform Limitations

Some of the current limitations of HP platforms include the following:

No Unified Extensible Firmware Interface (UEFI)

The UEFI interface is an industry standard for adding graphic management interface at the level of the basic input/output system (BIOS). This was originally developed by Intel Corporation, but is now managed by the Unified EFI Forum, as a standards body HP does have the ability to do remote management with its Integrated "Lights-Out" (iLO) "standard" for remote management capability. Unfortunately, iLO is only available on HP equipment, must be purchased with the server (advanced, iLO standard comes with most ProLiant servers), and must be paid for on an annual basis.

Fans don't have granular control

While HP does have system fans on both its C class and H class chassis, they have some significant differences from IBM fans. The HP fans cannot be regulated to the same degree of precision as IBM fans. Because there are more fans with less control, there will be more money spent on energy for the system fans, which will increase overall data center energy costs. In addition, the HP fans are significantly louder than IBM fans. While this is not likely to be a major issue in large data centers, since the trend is moving towards remote server management, this can certainly be a problem in smaller IT settings where servers are located in areas other than main data centers.

Dell Platform Limitations

Like the HP server product lines, Dell servers currently have no support for the UEFI standard. In a similar fashion, they do not support the same memory sizes, use less efficient power supplies (91% vs. 93%), and lack comprehensive integrated management. Dell also has 10u rack enclosures, which take up more space in the data center rack and their systems currently consume more energy than comparable IBM servers.

If the Dell M1000e enclosure uses six power supplies instead of three, it can be designed to be fully power redundant and dynamically reduce the energy consumed by power supplies on low usage, but this cannot be done if the system is not configured with redundancy. The Dell systems can be configured with nine redundant fans, which use "low flow" technology. This improves Dell air handling compared to its PowerEdge 1950 servers, but does not have as granular a setting capability as the IBM systems. As with HP, the Dell systems use a proprietary integrated Dell Remote Access Controller



(iDRAC) and CMC for its remote management interface. Memory is typically limited to 32GB (with plans for 64GB with eight GB DIMMS in the future¹².)

Conclusion

RFG believes power will be the number one issue for most large company IT executives to address for at least the next ten years, and probably even longer. IT managers will NOT have the option of ignoring this issue. One of the key ways IT executives can increase data centers efficiency is to use blade servers to increase the performance per watt of their systems. This reduces energy costs, while increasing infrastructure reliability and adaptability. The successful management of energy efficiency will be essential to saving money in the data center.

IT executives should incorporate power considerations in data center planning. This involves both working closely with facilities managers, as well as selecting systems that are going to work in a standards-based way with the environment. IBM can also help with this key issue by helping guide customers by using its power management tool¹³,

Since servers will occupy the greatest portion of data center space, with their density providing some of the greatest environmental challenges, RFG believes efficient processor and server design will be a critical factor for selecting data center servers. RFG believes Data center design and IT configuration, as well as the ability to virtualize the IT infrastructure are other key considerations optimize overall energy efficiency.

When looking at all these factors, RFG believes the combination of IBM BladeCenter systems with the new Intel® Xeon® 5500 Series processors provide the best combination of value, performance, and reliability. IBM has designed their blade servers with thoughts to today's requirements for environmentally friendly data centers. They have maximized the use of available memory, taking advantage of Intel Xeon 5500 Series on-processor memory management to improve memory performance. Intel and IBM have worked closely to optimize the use of dynamic power management to deliver the best performance per watt of each system, with the integration of UEFI improving remote management in a standardized way.

The constraints and demands of future data centers are more challenging than ever, with increased processor densities and increased budget limitations putting tremendous pressure on cooling, power, and space constraints. Blade servers can be an excellent option to increase the performance of business systems, but they will only deliver the best value where the processor and server architecture work synergistically to minimize limitations while maximizing capabilities.

¹² Current Dell PowerEdge Data Sheet Jan 2008

http://www.dell.com/downloads/global/products/pedge/en/pe_m_series_blades_specs.pdf

¹³ The power tool is available at <http://www->

[03.ibm.com/systems/bladecenter/resources/powerconfig/index.html](http://www-03.ibm.com/systems/bladecenter/resources/powerconfig/index.html)



120 Post Road West, Suite 201, Westport, CT 06880 Phone: 203-429-8950 Fax: 203-429-8930

Some systems, such as those from Dell and HP, have designs that address these issues, but not always in a synergistic fashion. HP uses more fans that consume more energy, and both HP and Dell use proprietary management interfaces that will not allow cross system integration. In HP's case, using all its iLO capabilities requires purchasing additional software licenses, increasing management costs. Cisco Systems has new blade systems coming onto the market, but its lack of experience in server design is likely to limit its ability to synergistically integrate architectural elements between the processor and server architectures.

IBM has the experience, reputation, and performance to effectively design and deliver integrated processor and systems designs for optimal blade systems performance in the data center that will meet future green data center requirements in a cost effective, reliable manner.

All rights reserved. 120 Post Road West, Suite 201, Westport, CT 06880. Telephone 203/429-8950. Facsimile 203/429-8930 www.rfgonline.com. This publication and all publications may not be reproduced in any form or by any electronic or mechanical means without prior written permission. The information and materials presented herein represent to the best of our knowledge true and accurate information as of date of publication. It nevertheless is being provided on an "as is" basis. Reprints are available for a nominal fee.

XSW03027-USEN-01