White Paper

WP01

Blade Server Technology Overview

October 2007

Blade.org

http://www.blade.org

I – Introduction to Blade Servers

The Blade server market has been the fastest growing segment of the worldwide server market for several years and the trend is expected to continue according to industry researchers. IDC reported in late August, 2007 that server blade market growth continued to accelerate with 36.7% year over year growth with customers increasing their blade deployments and vendors expanding their blade product portfolios. Why? IDC says "As IT organizations become more familiar with the platform, they are able to deploy blades in IT environments that are suited to take advantage of the management capabilities, as well as the cost and serviceability benefits."[1]

Blade servers offer a standardized method of deploying multiple processors, memory and I/O resources by placing those resources on plug-in boards that slide into a standard chassis. They reduce cost versus separate server systems by sharing common resources like power, storage and I/O across multiple processors. In the server market, the demand for blade servers has exploded as IT departments have found they are ideal vehicles for increasing server density and for gaining greater control over deployment, provisioning and upgrading large server farms. The combination of reduced cost, more efficient use of costly data center space, simplified management of tens to hundreds of servers, and greater performance have made blade servers irresistible.

Blade servers are both the beneficiaries and the drivers behind several key computing technologies, from increased processor densities in single-chip multi-core semiconductors, to cluster computing techniques, and 10Gb Ethernet deployments, as well as being intimately tied to advances in SAN and NAS storage systems, system management software and virtualization technologies. Blade servers ability to greatly increase both the quantity and the efficiency of computing, memory and I/O resources in a given physical space are a boon to data center managers who must satisfy ever increasing user demand for server productivity and has given rise to a new set of software tools both in server resource management and for server and desktop virtualization applications. In addition, the increasing densities of blade server-based data centers has given rise to a new set of challenges in the power and cooling of modern data centers that in turn are driving power vendors to develop new technologies and solutions for energy management and efficiency.

The key technologies involved in blade servers include:

- standardized chassis or enclosures with defined blade interface specifications
- multi-core processors
- standardized I/O blades and interfaces
- 10Gb Ethernet networking I/O
- Virtualization software for server consolidation
- Power and cooling monitoring and management

II – Blade Server Fundamentals

Blade servers start with a standardized chassis such as the BladeCenter H (see below) in which server blades are inserted. Each server blade is a separate system that runs its own operating system and application software independently of the rest of the system, relying on its local processor(s), memory and I/O resources to handle its applications. The server chassis contains the shared power, cooling and network connections to the rest of the data center.

Figure 1 – The IBM BladeCenter H provides 14 server blade slots in front and I/O module slots in the rear. Blade servers provide outstanding computing densities, high efficiency computing, low cost due to shared power, cooling and peripherals and outstanding flexibility.

As shown in the front and rear diagrams below[2], in the front there are power modules at the top and bottom of the chassis, with the main blade server slots arranged horizontally across the front of the enclosure. Media slots and USB connections are on the front right with indicator LEDs. In the rear, there are hot-swap blower modules, two management module bays, four standard fabric switch module bays, four high-speed fabric switch module bays, and a serial port breakout connector for direct serial connection.

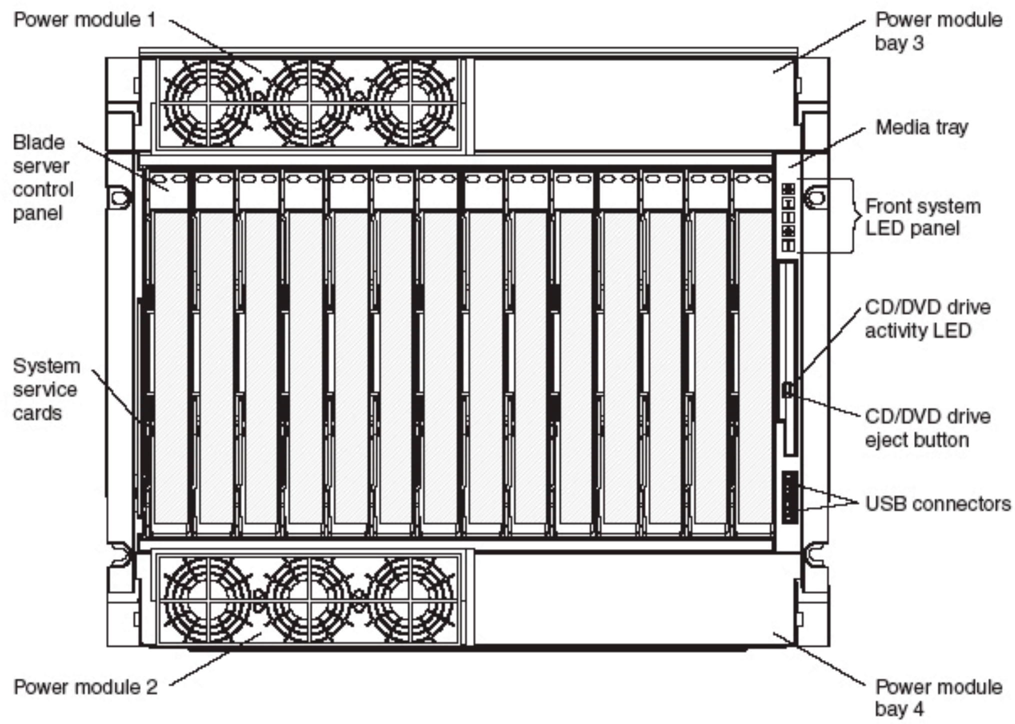

Figure 2 – The IBM BladeCenter H front panel diagram.

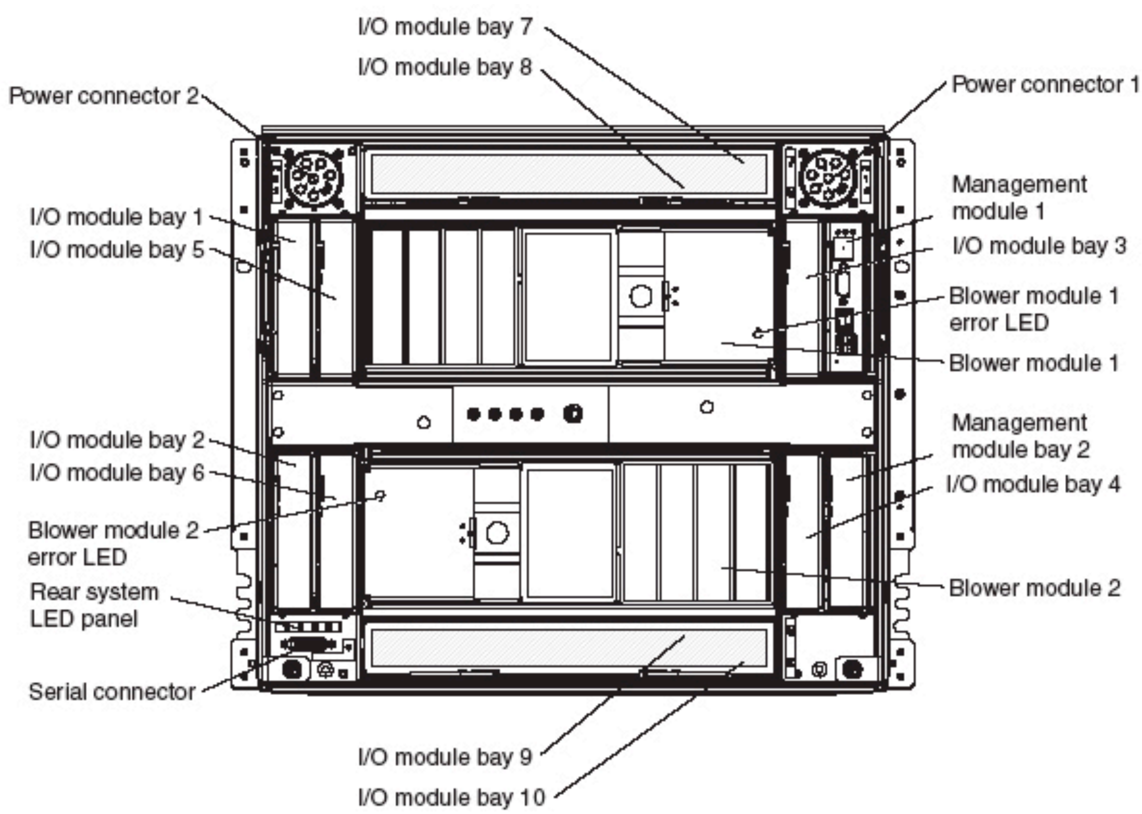

Figure 3 – The IBM BladeCenter H rear view diagram.

The BladeCenter H chassis supports either a 14 single-slot blade server configuration or a seven double-slot configuration. In each BladeCenter H chassis is a Management Module that manages the chassis and provides for KVM (keyboard, video and mouse) interfaces. A second management module can be installed for redundancy and backup. Users choose their I/O options such as Fibre Channel or Ethernet and install any desired I/O modules based on application requirements. The BladeCenter H chassis is in a standard rack form factor of 7U (height) and 28 inches (depth). Other chassis offer different sizes and differing abilities to support blade servers, I/O blades and network interfaces. Specialized enclosures such as the BladeCenter T chassis are optimized for particular applications, telecom applications in the case of the T chassis, or end user classes, such as the SMB chassis, intended for small to medium sized businesses.

With the blade enclosure containing all of the core shared computing resources, individual server blades hold a wide variety of computing, I/O and network options. Processor server blades support popular Intel and AMD single and multi-core processors with differing memory configurations. I/O blades provide various I/O options. Networking capabilities are obtained through switch modules that plug into I/O module bays and on individual blades through network interface expansion cards.

Processor Blades

Server blades provide the core processor, RAM and ROM memory, and interface to storage and network resources that are needed to provide server functionality. The server blade slides into the blade server chassis and connects to the shared resources, power, I/O, networking and storage through a backplane or midplane connection. The ability to add processing resources by simply sliding in another server blade is a key advantage that blade servers provide to the data center manager. Rack-based and standalone servers require much more installation overhead with individual cabling, power and networking infrastructure taking up considerable amounts of space and forcing a lot of manual interaction to setup and tear down. Blade servers, with their shared power and cooling, disk, CD and DVD media, switches and I/O ports, greatly simplify installation and maintenance in the data center.

A typical server blade is the BladeCenter HS20 from IBM. Two midplane connections link the server blade to the blade server enclosure. It is hot-swappable, allowing live upgrades to the server resource pool and comes in a variety of flavors featuring either low power operation or ultra-high performance extensions. There can be one or two processors on board, with from 1GB to 16 GB of main memory (RAM). The HS20 type 7981 for example, includes an Intel Xeon dual-core ultra low-power processor at speeds up to 2.0 GHz wth a front-side bus speed of 667 MHz, dual Gigabit Ethernet links, and onboard SAS storage controller with one or two internal hard disks drive holding up to 150GBs. The HS20 type 8843 on the other hand provides Xeon processors at up to 3.8 GHz with 64-bit extensions, a front-side bus speed of 800 MHz and onboard SCSI hard disks.

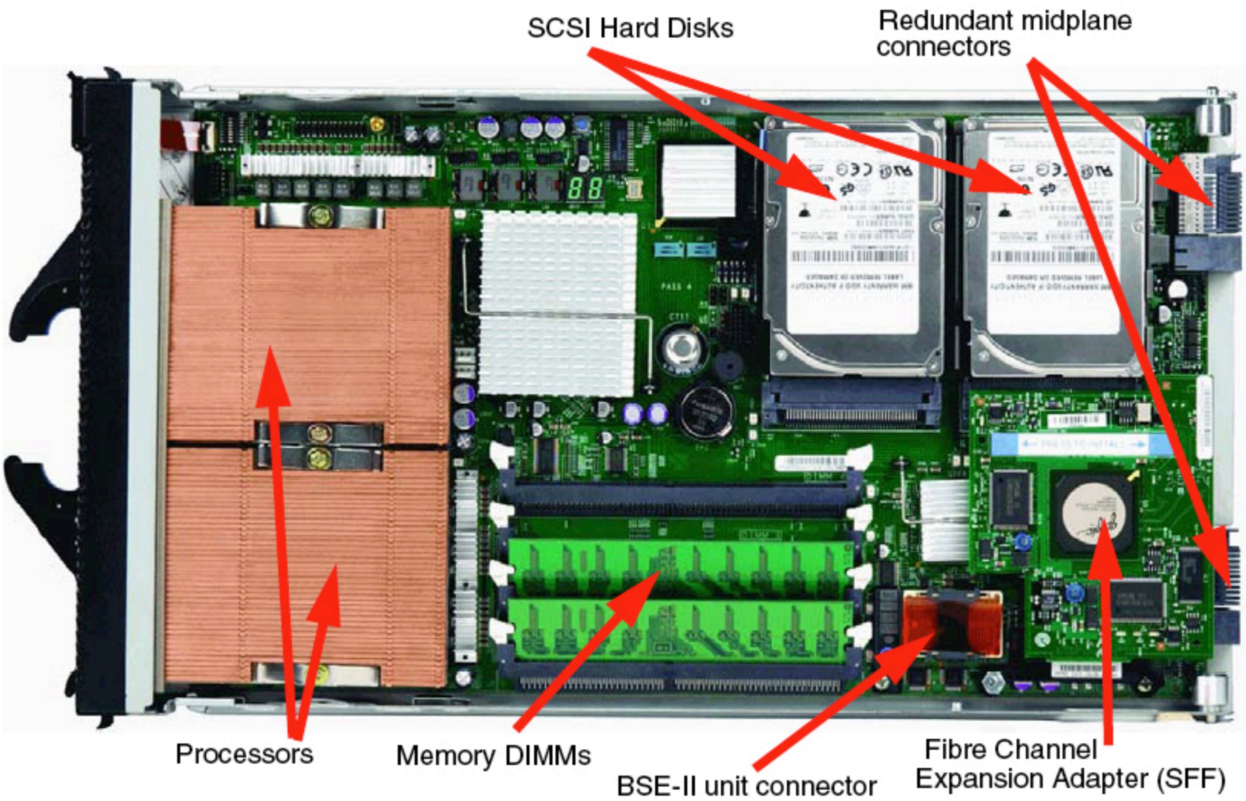

Figure 4 – The IBM BladeCenter HS20 server blade includes one or two processors, RAM memory, on-board hard drives and interface to the BladeCenter backplane to access power, I/O and network resources.

I/O Blade Modules

I/O blades provide specialized I/O and networking capabilities for BladeCenter platforms. A typical I/O blade is the Nortel Networks Layer 2-7 Gigabit Ethernet Switch Module, sold by Blade Network Technologies, that serves as a switching and routing fabric for the IBM BladeCenter Chassis.

Figure5 – The Nortel Layer2-7 GbE Switch Module blade includes four external 1000BASE-T connectors for up to 4 gigabit/sec. Ethernet connections, 14 internal full-duplex gigabit ports, one for each blade server in a BladeCenter chassis, management support for Spanning Tree Protocol and SNMP protocols and serial port for software installation and configuring.

One Nortel Gigabit ESM slides into a single I/O module slot shown on the right and left sides of the BladeCenter chassis in Figure 3 and can link all or any of internal 14 server blades in the chassis to a gigabit Ethernet network.

10Gb Ethernet Interfaces

With the increased densities now made possible in the data center with blade server technology, it has become increasingly beneficial to move to 10Gb Ethernet as a interconnect technology in the backbone of the data center infrastructure. New internal modules such as Nortel 10 Gb Uplink Ethernet Switch Module provide 10 Gb LAN switching and routing within the BladeCenter chassis and modules such as the Nortel 10 Gb Ethernet high speed Switch Module provide a full high-performance 10 Gigabit end-to-end connectivity from the blade server to the network. The advanced functionality of the Nortel 10 Gb Ethernet Switch Module requires a larger board, one that will fit only into the extra large I/O module bay slot 7 or 9 of a BladeCenter H or HT chassis.

Expansion Cards

While many I/O and network functions can be supported using plug-in I/O modules that fit into existing standard or extra large I/O module slots, it is useful or even necessary to have expansion cards that fit into the blade server modules themselves. The Nortel 10 Gb Ethernet high speed switch module, for example, requires that each connected server blade module that is connected to it be equipped with a 10 Gb Ethernet port, an option not installed in most server blades. Using an expansion card interface, this functionality can be added inexpensively to the server blade itself. For example the 10 Gb Ethernet expansion card from NetXen provides the necessary 10Gb Ethernet interface in an add-on card that plugs into a CFFh expansion slot.

Expansion cards fit one of five standard form factors as shown in the following table. The older blade server systems have a PCI-X connector while newer blade server systems have a PCI Express connector. The different expansion cards use either the PCI-X or the PCI Express connector to add their additional functionality. In the case of the Combination Cards, the CFFv can be used to extend to a CFFh. In the table, the outline of the expansion card is shown in red to show the relative positioning of the various expansion cards.

Expansion Card Form Factors			
Name	Abbreviation	Description	Photo (placement of card is shown in red outline)
Standard Form Factor	StFF	Standard Form Factor expansion cards is the original form factor for expansion cards in blade servers. An StFF expansion card uses the PCI-X connector. It use the same area on some blade servers as the internal disk drive.	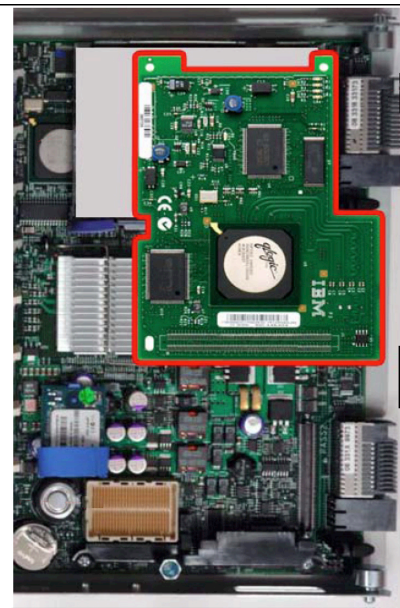
Small Form Factor	SFF	This card uses the PCI-X connector just like the StFF card, but it does not prevent the installation of a hard disk drive on the server blade module. Where the StFF card connects with the PCI-X connector and extends towards where the hard drive would be installed, the SFF card connects to the PCI-X connector and extends in the opposite direction.	

Combination Form Factor vertical	CFFv	The CFFv expansion card uses the PCI-X connector on the server blade or expansion unit and it extends the connecton to an additional CFFh or Combination Form Factor horizontal card.	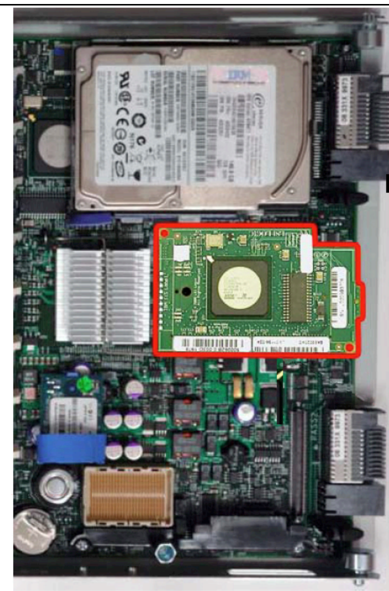
Combination Form Factor horizontal	CFFh	The CFFh expansion card connects to the PCI Express connector of the server blade and can be used in conjunction with a CFFv adapter card.	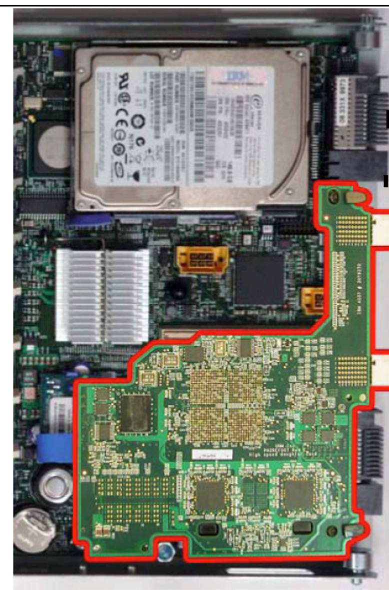

High Speed Form Factor	HSFF	The HSFF expansion card connects to the PCI Express connector and when installed, no other card can be installed on the PCI-X connector. Only one HSFF adapter can be installed on a single server blade.	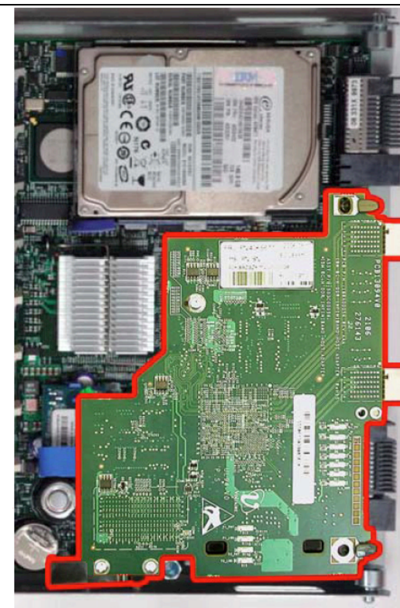

The NetXen 10 Gb Ethernet card is available in the combination form factor horizontal or CFFh format as shown in the figure below. The NetXen network interface card, called an Intelligent NIC, is based on a highly integrated, fully programmable network controller solution that integrates dual 10 Gigabit links with native 4 GB/s x8 PCI Express interfaces. The Intelligent NIC can support a wide variety of protocols including TCP/IP, RDMA, iSCSI and security.

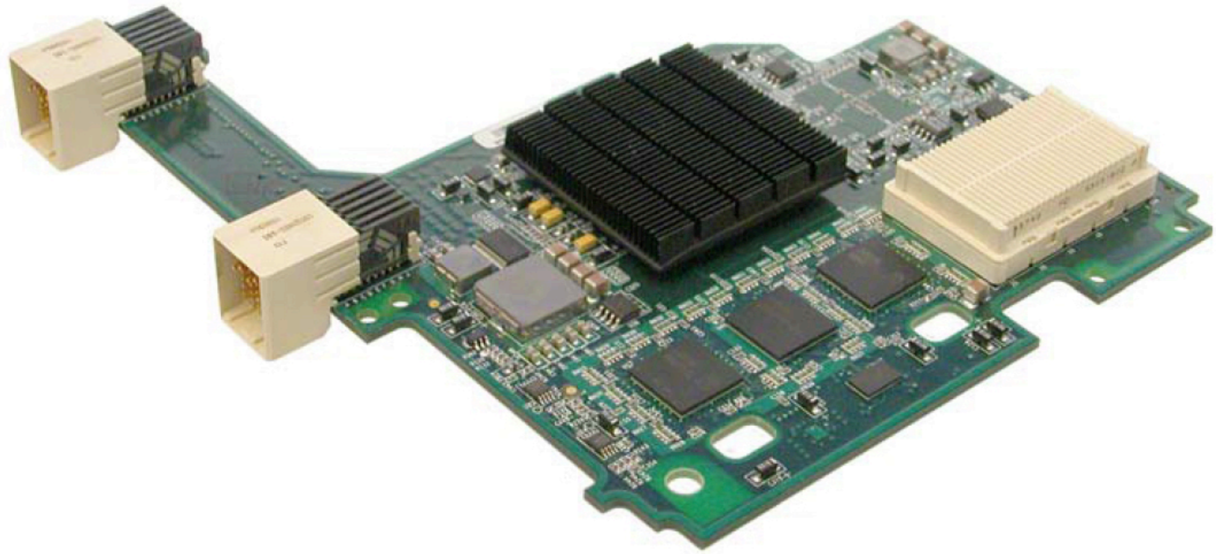

Figure 6 – The NetXen 10 Gb Ethernet Expansion Card is available in the CFFh form factor. It includes dual 10 Gb Ethernet interfaces and an x8 PCI Express interface and uses the NetXen NX2031 Intelligent NIC network controller chip.

III – 10Gb Ethernet Network Technology Emerging in Blade Servers

Blade servers are accelerating the deployment of 10 GbE networking in the data center infrastructure because the high bandwidth links are well suited for connecting dense computing resources such as blade servers to each other and to extended data center resources. Ethernet is a well known and easily supported technology and 10 GbE fits with comfortably in existing 1 Gb Ethernet installations. Recent developments such as Fibre Channel over Ethernet and the growing popularity of iSCSI for storage networking have made 10 GbE the logical choice for many to become a unifying networking protocol for heterogeneous environments such as a high performance data center where infrastructure may include Ethernet, Fibre Channel, iSCSI and InfiniBand among others.

A key driver in the demand for the extra bandwidth delivered by 10 GbE links is the trend towards server virtualization. Blade servers can host 10-14 physical server blades, each with up to 2 to 8 processor cores. With virtualization, from 10 to 100 virtual machines can be created on each blade, yielding a compact blade server system with hundreds, if not thousands of virtual servers accessing the data center infrastructure. One of the virtues of blade servers is the ability to obtain greater efficiency and computing density. But with this increased server density comes a correspondingly greater demand on the communication links between individual physical servers, between servers and network attached storage and storage area networks and between servers and the WAN network. 10 Gb Ethernet links can handle the increased aggregate bandwidth generated by blade server traffic. In addition, the ability to support popular storage based protocols such as Fibre Channel (in Fibre Channel over Ethernet) and iSCSI enables 10 Gb Ethernet to be a singular pipe that is easier to manage, scale and maintain, while supporting the underlying heterogenous protocols already existing in the data center.

With the concentration of resources in blade server environments, disaster management becomes a more important issue, as well. The high bandwidth of 10 Gb Ethernet links support a quick backup and restore capability in the blade server environment.

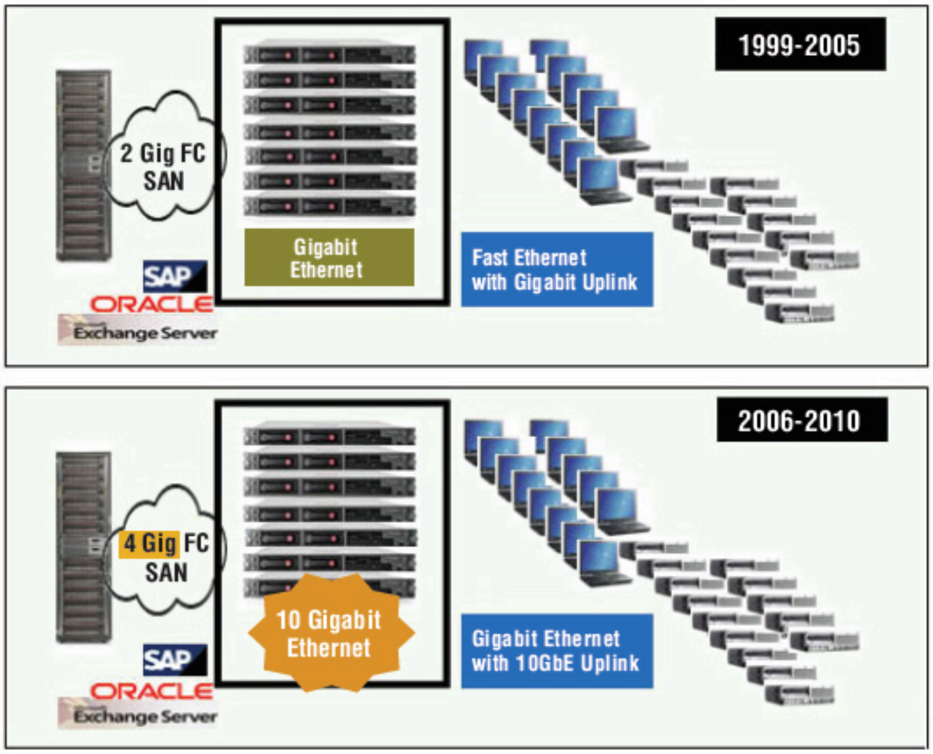

* Source: NetXen, Inc.

Figure 7 – The network infrastructure is seen as moving from 100 Mb/s Fast Ethernet with 1 GbE uplinks to a GbE infrastructure with 10Gb Ethernet uplinks to blade server environments using 10 Gb Ethernet internally and externally to carry 2 and 4 Gb Fibre Channel traffic from Storage Area Networks.

The appeal of 10 Gb Ethernet in the data center goes beyond the raw physical increase in bandwidth. As blade server processors are occupied with tens or hundreds of virtual machines, their OSs and the applications that run on them, they have less processor bandwidth to handle the normal network processing chores of TCP/IP and Ethernet packet processing. 10 GbE solutions often have hardware acceleration and TCP/IP packet processing offload built into the 10 GbE NIC. This greatly reduces the CPU utilization absorbed by network processing overhead, freeing up the processor for more server processing.

Server consolidation, server virtualization, storage consolidation, and achieving an unified and integrated networking infrastructure are all trends and goals that are accelerating the deployment of 10 Gb Ethernet networking technology in the blade server environment.

IV – Virtualization

As previously noted, virtualization has become an important blade server technology because it is a means of achieving server consolidation, increasing server utilization and efficiency, simplifying data center management and improving performance. So, what exactly is it?

Server virtualization is an abstraction technology that enables the division of the hardware resources of a given server into multiple execution environments and enables the consolidation of multiple servers and hardware resources into a single computing resource. Essentially, virtualization software is a layer of software that runs between the hardware of the server and the OSs and applications that run on the server. It controls access to the server hardware resources and creates a software image of the real hardware system and presents that image to the OS and applications. Since it is just another software program, it can create multiple instances of “virtual machines” and have OSs and applications run on those instances. There are two parts to the virtualization software. The hypervisor or Virtual Machine Monitor creates and manages the second part, the virtual machines that mimic the operation of a real physical system. On top of these virtual machines, standard OSs and standard applications run normally, unaware that they are executing on a virtual rather than a real computing system.

Virtual machines created through virtualization software can:

- consolidate the workloads of several servers into one, higher performance server
- create the illusion of legacy hardware so that investments in existing software/hardware infrastructures can be maintained even as a data center migrates to newer, more productive and higher performance hardware configurations
- be used to run multiple operating systems and applications simultaneously, increasing the flexibility and utility of a given set of data center resources

- enable the isolation of applications within a single server/server complex so that runaway applications do not bring down the entire server infrastructure
- give data center managers a simplified infrastructure management control point that gives better control over the common tasks of server provisioning, backups and disaster recovery, and system migration.

Virtualization has emerged as key technology for servers and blade servers because of the many benefits it provides to the data center manager. Virtualization gives the IT manager the ability to see and manage computing resources in a more flexible, more efficient and much more cost effective way. Virtualization provides a means for increasing server utilization, implementing server consolidation, increasing flexibility by supporting multiple and legacy OSs and applications in a modern hardware environment and improving data center management.

By implementing a software virtualization hypervisor and running multiple instances of virtual machines, a single server or a pool of integrated servers can provide a resource pool equivalent to many times more individual, separate server systems. This greatly lowers the cost of data center deployments and gives data center managers superior tools for rapid deployment of applications, better disaster security and lowers the total cost of ownership. While initially a software-based technology, advancements in integrated virtualization support features by server processor vendors have made virtualization a software and hardware technology with the benefit of even greater virtual machine performance.

The latest versions of server processors from AMD and Intel include special features and instructions that are designed to optimized and streamline the creation and management of virtual machines within the x86 architecture. These newer “virtualization-enabled” processors add hardware acceleration to many common virtualization tasks such as switching from one virtual machine to another and greatly simplify the task of creating optimized virtualization hypervisors.

V – Power and Cooling Challenges

Of course, the consolidating servers, increasing server efficiency, packing greater computing power into smaller, more compact physical spaces puts a greater strain on the powering and cooling of the datacenter. The advantages gained by moving to a blade server environment opens up the opportunity to put more server compute resources into the same physical data center space. Data centers designed for one server/one hardware platform architectures may become overwhelmed with the power and cooling requirements of tightly packed blade servers running at near peak performance levels 24x7. To compound matters, the cost of energy has skyrocketed.

As shown in the figure, analysts predict that the cost to power and cool data centers will become greater than 50% of the cost of the equipment itself.

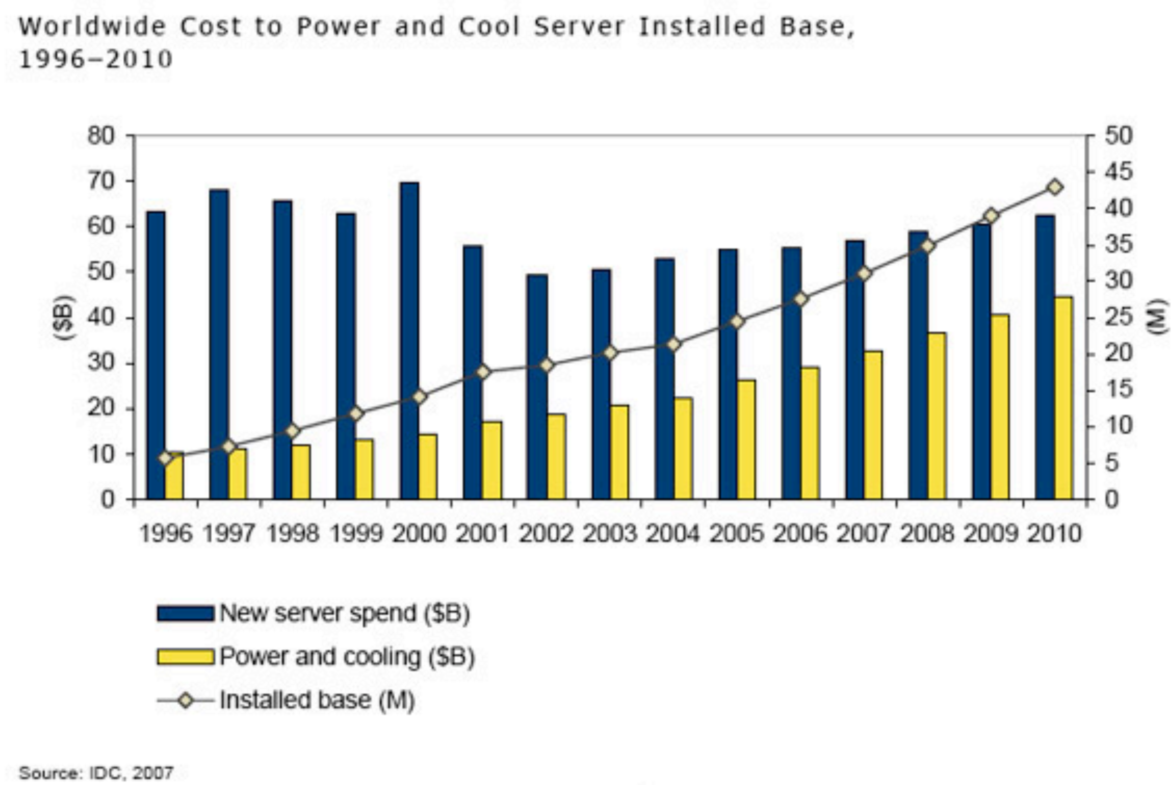

Figure 8 – Market researcher IDC projects the expense to power and cool the installed base of servers will increase at four times the growth rate for new server spending.

To cope with these spiraling power and cooling costs, data center manager must take steps to plan for present and future power consumption. If best practices are achieved, dramatic energy savings can be realized. As shown the chart below, the EPA estimates that depending on the deployment of energy efficient technologies and practices, energy consumption in the datacenter can be reduced from the historical high year to year growth rates to flattened or even reduced energy consumption – even with the performance of data centers continuing to increase at their historical pace.

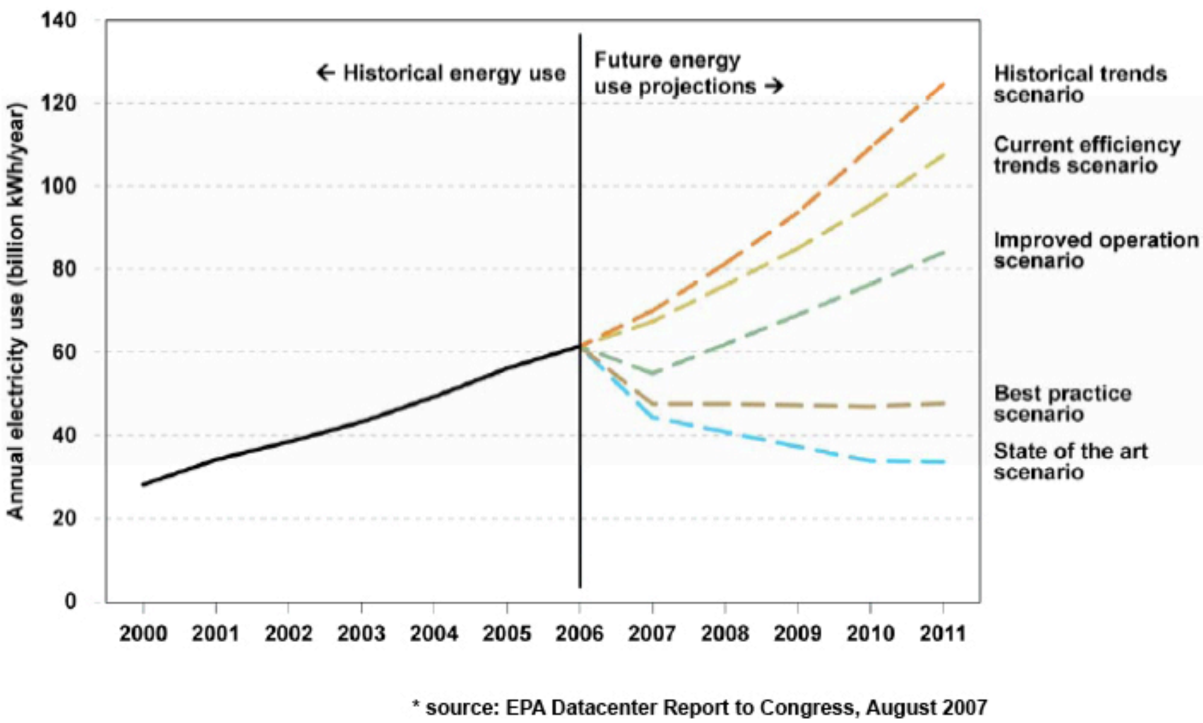

Figure 9 – The EPA report to Congress predicts three distinct energy consumption options based on how the industry adopts energy efficiency technologies and techniques. Worst case, no change in the increasing energy consumption and costs. Best case, there are significant savings in energy consumption and costs.

According to the EPA report, improved operational technologies can lead to savings of 23 B kWh by 2011. Best practice operation of data centers can lead to 60 B kWh and state-of-the-art energy efficiency deployments could lead to 74 B kWh of electricity consumption savings in 2011. This leads directly to cost savings ranging from $1.6 B to over $5.0 Billion.

Current energy efficiency trends that are already underway are expected to yield only minor results over the next five years. Therefore, it is obvious that more must be done to reduce the cost of power and cooling for the datacenter. For the entire data center, vendors such as Blade.org members AMC-MGE and Eaton have comprehensive solutions ranging from power supplies and cooling systems to entire integrated data center systems to power and cooling consulting services. Individual vendors ranging from processor vendors creating lower power processors, networking vendors creating lower-power I/O subsystems, and storage subsystem makers producing lower power mass storage systems, are contributing to the more efficient data center. For blade servers specifically, there are both individual products and management tools that can help the data center manager monitor and manage power consumption, the first step in understanding and controlling runaway energy consumption and cost.

In the EPA report, the different scenarios, ranging from improved operation, to best practice, to state-of-the-art, vary in the expected deployment of IT equipment and site infrastructure system deployments. For example, in the improved operation scenario, it is assumed that virtualization technology can lead to a physical server reduction ratio of between 1.04:1 to 1.08:1 by 2011 and energy efficient servers will grow to 15% of deployments with power management systems enabled in 100% of servers. In the state-of-the-art scenario, virtualization ratios go up to 1.66:1 to 5.00:1 by 2011 with energy efficient servers representing 100% of volume server shipments from 2007 to 2011.

The key differences in site infrastructure scenarios were in deploying more efficient transformers (95% to 98% efficiency transformers), increased use of efficient un-interruptable power supplies or UPS (from 80% to 95%), switching from air cooled direct exchange system chillers to liquid cooling with cooling towers, and deploying variable-speed fans and pumps.

Specifically for blade server environments, one of the key factors can be the increased use of virtualization to enable server consolidation and reductions in networking and storage hardware and the use of improved power monitoring and control software management tools.

For example, as shown in the figure below, only a fraction of the expensive electrical energy input to a datacenter is used by the IT equipment. The rest is consumed in non-compute sources.

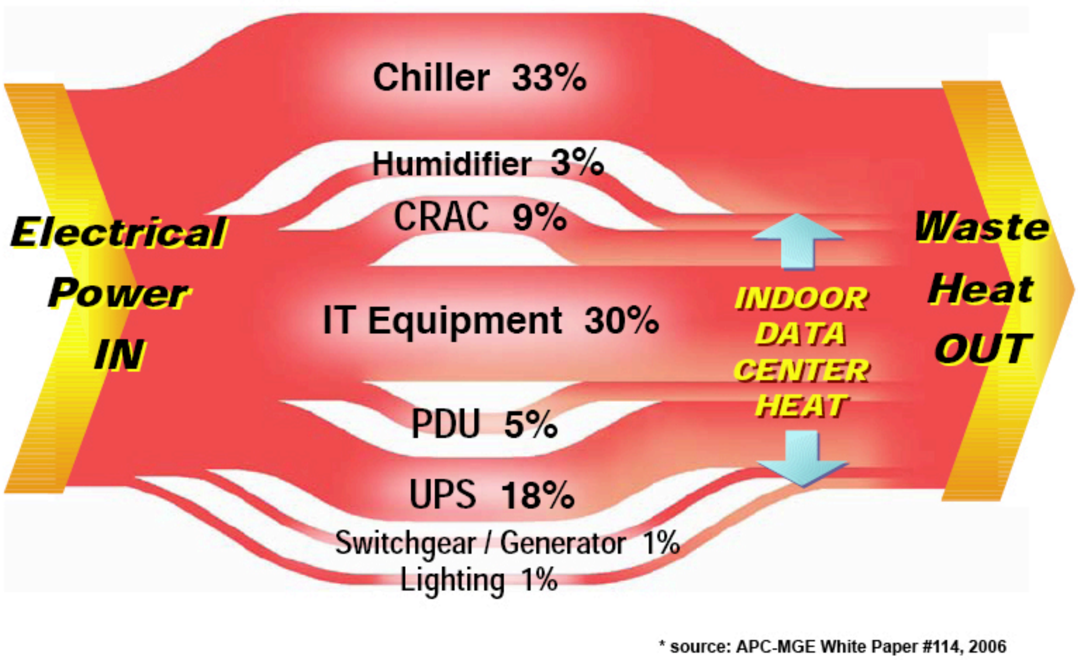

Figure 10 – The energy into a datacenter gets wasted if it is not used on the IT equipment itself. Current analyses show that to be the case today.

Monitoring and Managing Power Consumption

Of course, to manage power consumption, one must employ advanced power-monitoring software to understand and control power consumption. It is not uncommon for data center managers to lack information on how much power individual servers consume. Without an accurate way to measure power consumption, peak loading and the effects of virtualization, adding more servers, switches or storage subsystems, it is impossible for managers to effectively manage power consumption. New advanced tools such as IBM PowerExecutive provide automated tools needed to monitor and manage power consumption accurately. It can measure real-time power consumption and heat emission by individual server, server group or location. It allows the optimization of energy use and the lowering of power consumption when low utilization can provide cost savings. These power monitoring and management capabilities are an important tool in achieving energy efficiency in the data center.

One of the key advantages of PowerExecutive is that it measures actual power usage across multiple servers. Information such as power consumption, thermal information regarding ambient temperatures and BTUs generated by an individual, or group of blade or rack servers, and for BladeCenter H installations, exhaust temperature can be captured and displayed. Armed with this information, a datacenter manager can properly size power input requirements and more accurately plan datacenter power requirements.

Here is an effective strategy for reducing power consumption today:

- Make a plan - start with assumptions, if necessary and then fine tune as more data about the data center consumption profile is obtained
- Monitor and then analyze data center energy consumption and factor those results back into the plan
- Develop a timeline and forecast for future data center performance requirements
- Explore which new technologies, blades, virtualization, infrastructure consolidation, power management programs are available to manage and reduce power consumption
- Use that plan with expert assistance from energy and equipment vendors to map out a realistic program of data center floor planning, cooling planning and power management
- Continue to monitor and analysis power usage and use that data to refine data center planning

APPENDIX 1 – BladeCenter Chassis Models

BladeCenter Chassis			
Chassis Name	**Description**	**Size**	**Extras Supported**
IBM BladeCenter E 	Provides the greatest density and common fabric support and is the lowest cost option	H - 12.0” W - 17.5” D - 28.0”	
IBM BladeCenter H 	Delivers the high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments	H - 15.75” W - 17.4” D - 28.0”	High-speed I/O bays for high-speed switches
IBM BladeCenter T 	Designed for telecommunications network infrastructures and other rugged enviroments	H - 13.75” W - 17.4” D - 20.0”	
IBM BladeCenter HT 	Designed for high-performance, flexible telecommunciations environments with support for high-speed internetworking technologies such as 10G Ethernet.	H - 21.0” W - 17.4” D - 27.8”	High-speed I/O bays for high-speed switches

APPENDIX 2 – BladeCenter Blade Server Modules and I/O Blades

Blade Server Modules for BladeCenter			
Server Blades			
Blade Name	**Description**	**Size/ Form Factor**	**Chassis Supported**
H20	Server blade with one or two Intel Xeon processors, up to 16GB RAM memory, I/O and internal HDD storage	1 slot	H, E, T, HT
HS21	Server blade with one or two Intel Xeon multi-core processors (up to 8 cores) with up to 32GB RAM memory, on HDD drive, 2 GbE ports, plus I/O and network expansion	1 slot	H, E, T, HT
HS40		2 slots	H, E, T
LS20		1 slot	H, E, T, HT
LS21		1 slot	H, E, T, HT
LS41		1 or 2 slots	H, E, T, HT
JS20		1 slot	H, E, T, HT
HC10		1 slot	E
I/O Modules			
Cisco Systems Fiber Intelligent Gigabit ESM	Ethernet Switch Module	1 slot	E, H, T, HT
Cisco Systems Intelligent Gigabit ESM	Ethernet Switch Module	1 slot	E, H, T, HT
IBM Server Connectivity Module	Ethernet Switch Module	1 slot	E, H, T, HT
Nortel L2/3 Copper GbE Switch Module	Ethernet Switch Module	1 slot	E, H, T, HT
Nortel L2/3 Fiber GbE	Ethernet Switch Module	1 slot	E, H, T, HT

Switch Module			
Nortel Layer 2-7 GbE Switch Module	Ethernet Switch Module	1 I/O slot	E, H, T, HT
Nortel Layer 2/3 10 Gigabit Uplink Ethernet Switch Module	Ethernet Switch Module	1 I/O slot	E, H, T, HT
Nortel 10 Gb High Speed Switch Module	Ethernet Switch Module	1 I/O slot	H, HT
Ciso Topspin InfiniBand Swtich Module	InfiniBand Module	1 I/O slot	E, H, T, HT
Cisco Systems 4X InfiniBand Switch Module	InfiniBand Module	1 I/O slot	H, HT
QLogic InfiniBand Ethernet Bridge Module	InfiniBand Module	1 I/O slot	H, HT
QLogic InfiniBand Fibre Channel Bridge Module	InfiniBand Module	1 I/O slot	H, HT
Brocade 4 Gb 10-port SAN Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
Brocade 4 Gb 20-port SAN Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
Cisco Systems 4 Gb 10-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
Cisco Systems 4 Gb 20-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
QLogic 4 Gib 10-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
QLogic 4 Gib 20-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
QLogic 10-port 4 Gb SAN Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
QLogic 20-port 4 Gb SAN Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
QLogic 4 Gb Intelligent Pass-thru Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT

McDATA 4 Gb 10-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
McDATA 4 Gb 20-port Fibre Channel Switch Module	Fibre Channel Switch Module	1 I/O slot	E, H, T, HT
IBM BladeCenter Optical Pass-thru Module	Pass-through and Interconnect Module	1 I/O slot	E, H, T, HT
IBM BladeCenter Copper Pass-thru Module	Pass-through and Interconnect Module	1 I/O slot	E, H, T, HT
IBM Multi-Switch Interconnect Module	Pass-through and Interconnect Module	1 I/O slot	E, H, T, HT
Expansion Cards			
BladeCenter Fibre Channel Card	Expansion Card	StFF Form factor	Plugs into a variety of server blades
BladeCenter 2 Gb Fibre Channel Card	Expansion Card	SFF	"
BladeCenter Gigabit Ethernet Card	Expansion Card	StFF	"
BladeCenter Gigabit SFF Ethernet Card	Expansion Card	SFF	"
Cisco Systems InfiniBand 1X HCA Card	Expansion Card	StFF	"
Cisco Systems InfiniBand 4X HCA Card	Expansion Card	HSFF	"
Cisco Topspin InfiniBand HCA	Expansion Card	StFF	"
Emulex 4 Gb SFF Fibre Channel Card	Expansion Card	SFF	"
Myrinet Cluster Expansion Card	Expansion Card	StFF	"
QLogic iSCSI Expansion card	Expansion Card	StFF	"
QLogic 4 Gb SFF Fibre	Expansion Card	SFF	"

Channel Card			
QLogic 4 Gb StFF Fibre Channel Card	Expansion Card	StFF	"
IBM Concurrent KVM Feature Card	Expansion Card	--	"
NetXen 10 Gb Ethernet Card	Expansion Card	CFFh	"
Ethernet Expansion Card	Expansion Card	CFFv	"
QLogic Ethernet and 4Gb FC Card	Expansion Card	CFFh	"
QLogic 4 Gb Fibre Channel Card	Expansion Card	CFFv	"

End of Blade.org WP_01 White Paper.

Notes:

1 – IDC Worldwide Quarterly Server Tracker, August 23, 2007.

2 – IBM BladeCenter Products and Technology Redbook SG24-7523-00, August, 2007.

3 – IDC White Paper, Solutions for the Datacenter's Thermal Challenges, January 2007.

4 – EPA Datacenter Report to Congress, August 2, 2007.

About Blade.org

Blade.org is a collaborative organization and developer community focused on accelerating the development and adoption of open blade server platforms. The organization was established in February 2006 to increase the number of blade platform solutions available for customers and to accelerate the process of bringing them to market. From eight founding companies, Blade.org has grown to nearly 100 members including leading blade hardware and software providers, developers, distribution partners and end users from around the globe.

For more information, please visit: http://www.blade.org.