White Paper

FOCUS consulting

HP Blades and Server Virtualization:

The Ins and Outs of I/O

Sponsored by

CDW®

hp® invent

The combination of blade servers and virtualization today offers a compelling solution for IT departments looking to regain control of their server resources, while offering benefits including space savings, reduced power and cooling, improved manageability, availability and flexibility. Initially, blade servers got a bad reputation for not being a good match for virtualization. This was generally because early blade servers were not powerful enough, and often did not have enough Input/Output (I/O) resources - Ethernet NICs and Fibre Channel ports - to support multiple virtual server environments on one blade, not to mention hundreds of virtual servers in a single chassis. As with all technology, both blades and virtualization software have evolved to better support the high performance and high bandwidth I/O requirements that come with initiating many virtual machines.

In 2001, when server blades first emerged, they were designed to target a very narrow part of the server market — specifically scale-out applications, such as large web server farms, where the footprint, or the amount of floor space required, was the biggest concern. Immediately, IT liked the "idea" of server blades with shared power, cooling, I/O, and management, but, they wanted the blades to have comparable capabilities to their rack mounted server options, such as processing power, memory architectures, and I/O connectivity. In other words, although these first blades were not meant to be general-purpose servers, IT users wanted them to be. Based on user input, many blade server vendors responded releasing new blade servers that meet and even exceed the capabilities found in rack servers.

> "…blades were easier to install, easier to stock unified spares for increased availability, and easier to swap with more predictability in the amount of time needed to swap…"

Virtualization has had a similar history, in that originally, it was used in a fairly narrow market - the mainframe/large computing market - primarily for development and testing. This specialized software allowed IT to efficiently utilize these expensive computer resources to their utmost capacity. Again, around 2000, a startup company called VMware brought virtualization technology to industry-standard x86 computers. Their first server virtualization release was in 2001 with VMware GSX Server. Today, they have improved upon and expanded their offerings to the most recent and sophisticated release of VMware ESX, Virtual Infrastructure 3 (VI3).

Together, server blades and virtualization are addressing many of IT's cries. "My servers run at 3 — 5% utilization rates and I'm being asked to provision another ten servers each running their own application. I'm out of space; I'm out of power and I can't stretch my budget or my administrators any thinner. How can this be good for business?" However, the question arises, "Can the system architectures, I/O and operating environments designed for a server running one application translate to support multiple virtual servers and their applications?"

Today, the answer to this question is a resounding "Yes!" for many application environments and workloads. As with any technology, early blades and server virtualization solutions experienced growing pains, both real and rumored. In particular, I/O throughput, provisioning, and ongoing management presented challenges in the beginning. Over the past two years, chip, blade and server virtualization vendors made important changes to improve data movement into and out of the servers. This paper discusses the architectural and operating environment improvements that allow for faster data handling in server blades and server virtualization technologies. It will provide tips on how to determine if blades and virtualization can meet or exceed your business needs.

Virtualization on Blades — an architectural perspective

In an effort to address server sprawl and consolidate the plethora of distributed, underutilized servers, many businesses are turning to server virtualization technology on blade servers. These technologies together can reduce the amount of physical space, power and cooling and costs required for servers significantly. When these two technologies are combined, IT also gains increased server utilization rates, higher reliability, flexibility, and serviceability. To get the most out of these technologies, it is important to understand how the solution architectures can help or hinder these goals.

Server Virtualization Architecture

With increasing cost pressure on IT organizations and underutilized servers taking up space, power and administrative cycles, not to mention putting out heat, it's not surprising that many SMB and enterprise organizations have been looking for solutions. Server virtualization has allowed companies to consolidate anywhere from 4:1 to 20:1 virtual servers to physical servers for many of their application environments.

VMware was the first to deliver virtualization to industry-standard x86 servers. Their initial server product, now called VMware Server, is based on a virtualization architecture, commonly called "hosted," as shown in Figure 1, where each virtual server talks through the native operating system to access devices.

Hosted Virtualization Architecture

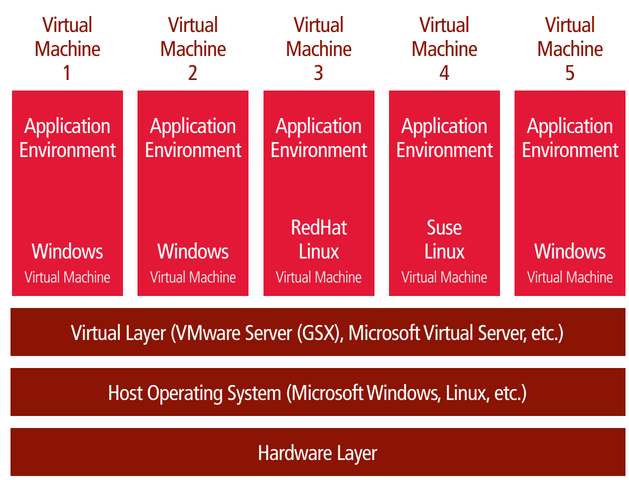

Figure 1

This architecture, while fairly easy to develop and implement, has significant overhead (performance) issues. As shown in the diagram, each application sits on top of an unmodified operating system which sits on top of a virtualization layer which sits on top of another operating system (which may or may not have been modified to run as a virtual host server). This creates many layers for I/O or other hardware-specific commands to move through before reaching the hardware. Each layer adds a little bit of overhead which adds up to slower performance. For users who would like to test virtualization or who don't have stringent performance requirements, VMware Server, which is available for free, can be a good option.

> "With virtual machines, IT no longer has to go through the long cycle of justifying new [physical] servers every time a new application is introduced; we simply add another virtual machine… Our BladeSystem puts us in a better position in terms of availability — if something fails, the VM's can be moved to another blade. For example, during the implementation, a loop failed on the SAN. The VMs lost all connectivity for a second, but the servers stayed online. When the loop failed over, all went fine with no staff intervention and the VMs and their applications continued… Today, our blade servers run between 24% and 70% utilization, depending on the application mix."

For those with higher performance requirements, VMware and other vendors have developed a thin hypervisor that is booted instead of an underlying operating system. The hypervisor then creates virtual machines which contain the native operating systems (see Figure 2). The overhead reduction when using a hypervisor architecture is fairly obvious over the hosted architecture.

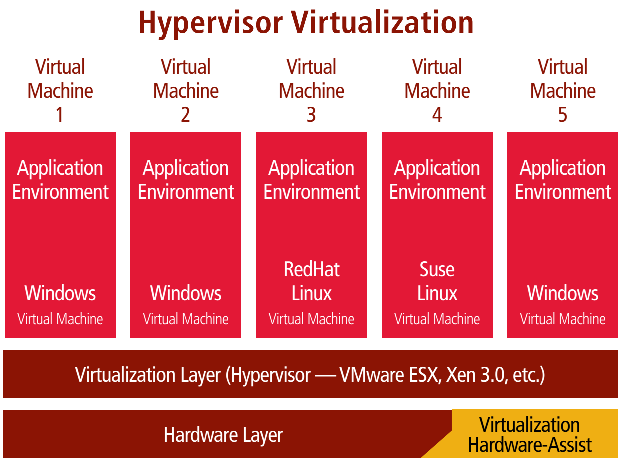

Figure 2

The number of virtual machines and associated operating and application environments one can run on any given server will depend on the server configuration as well as the requirements of the distinct applications. This variability is one reason why blade solutions fit so well with virtualization solutions. Companies have a number of applications with differing processor, memory, storage and networking options and the modularity of blades allows each blade server and blade chassis to be configured with differing processor, memory, storage and networking components to fit the requirements of the virtualized environments, while maintaining a level of standardization.

Running multiple virtual servers works the same way on a blade server as on any other server. First, a hypervisor is booted on the server that contains the base functionality needed to interact with the devices and manage the virtual machines. From there, VMware provides management software called VirtualCenter in which virtual machines are configured and initiated. It is important that the chosen virtualization solution provides adequate management and monitoring capabilities to easily track the performance of applications within the virtual machines on any given blade server. VMware, as the most mature product in the marketplace, has the most extensive management suite, including the capability to perform live migrations or movement of running VMs from one blade to another, VM failover from one server to another, and non-disruptive backup of the VM environment. VMware tools also support Dynamic Resource Scheduling (VMware DRS) which will continuously monitor the physical server loads against given thresholds and move VMs around to maximize utilization across all the physical servers.

Data Storage and Virtualization

In order to support the ability to quickly and easily move VMs from one server blade to another, VMware creates a VMware ESX Server File System (VMFS) that can be accessed by all the virtual servers (ESX servers) within a resource group as shown in Figure 3. The VMFS is a clustered file system that enables each of the VMs to have access to their data no matter which ESX server the VM is running on. The VM operating environment is given a virtual LUN (logical unit number) to send and receive SCSI commands and data. The ESX servers work in conjunction to provide data locking and access restrictions in order to prevent data from being accessed or overwritten by mistake.

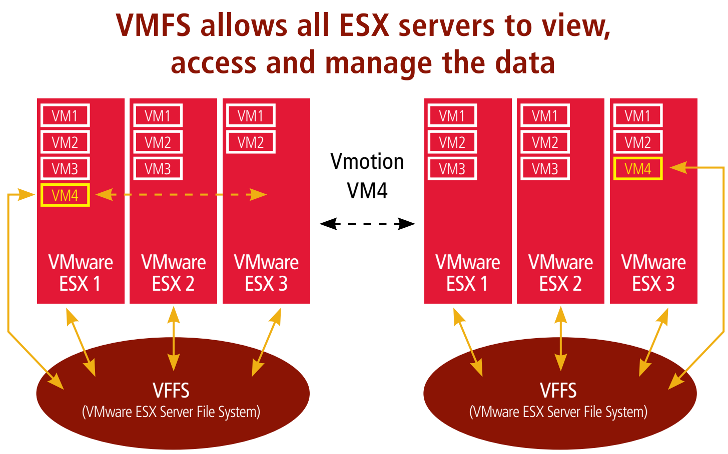

With VMFS, a Virtual Machine (VM) can access its data from any ESX Server sharing the same VMFS

Figure 3

This high performance file system can span across the disks within a SAN (Fibre Channel or iSCSI) or on NAS appliances allowing all VMs access to their data regardless of which ESX server they are running on. It is also possible to directly attach a SAN LUN to a virtual machine (discussed later in the *Migrating from Physical to Virtual* section).

Blade Architecture

Much of the value of server blade systems comes from the shared components within the blade chassis, making them an excellent partner for virtualization — further consolidating server space and sharing components. In a blade chassis, the backplane, the power supplies, the cooling/fans, the I/O modules and the management, are shared by all the server blades held in the chassis. For example, the HP BladeSystem c7000 chassis holds up to sixteen half-height or eight full-height blades in 10U of space; with the blades sharing the rest of the chassis components (see Figure 4).

BladeSystem Shared Architecture Cross-Section

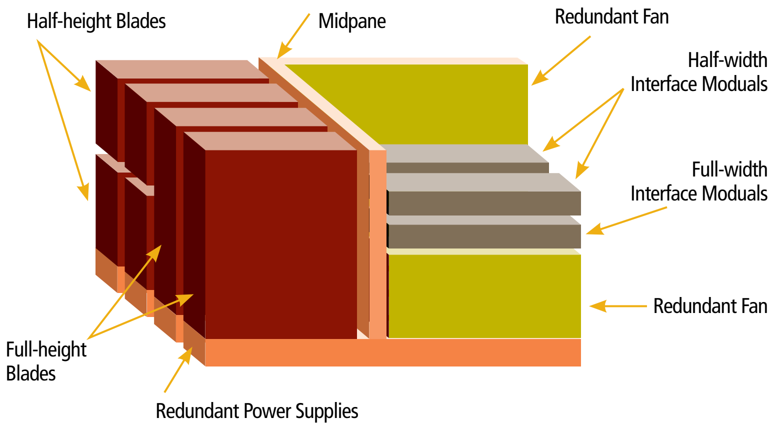

Figure 4

The redundant, shared components ensure high reliability, and both the blade and interconnect options provide increased flexibility of configuration. The fact that all these components are hot-swappable creates a highly serviceable solution.

I/O Options

The combinations of options for storing and moving data through a blade are numerous. Each blade can be configured with or without local on-blade storage, as well as with various alternatives for connecting to networked storage (SAN or NAS). In addition, HP users have the option of adding local storage through storage blades. For SAN connections, mezzanine cards can be added for additional IP network ports, Fibre Channel SAN ports, or InfiniBand connections. On the other side of the midplane, Ethernet, Fibre Channel, and InfiniBand switches can be incorporated or Pass-Thru modules can be used. HP is also now making great strides in virtualizing I/O within the HP BladeSystem, with its new Virtual Connect technology.

Blade Storage Options

Most server blades have the option to be configured with on-blade storage or without it (called diskless.) The HP c-Class blades support two or four (depending on the blade) small form factor (SFF) disks with a hardware RAID controller for reliability and speed. The disks drives are located at the front of the blade. This gives users easy access and hot-swap capabilities in the event of a disk failure. Simply pop the bad disk out and pop a new one in, without ever removing the blade from the chassis or stopping your application processing. It's important to note that not all vendors support hardware RAID on their server blades, nor are all blades architected to support drive hot-swap capabilities.

While still not the standard configuration, the idea of having diskless server blades is becoming more popular. There are two primary drivers for removing local storage. The first is that the disk is the only moving component on the blade, which makes it more likely to fail. If you remove the disks, the mean time between failures (MBTF) on a blade is increased. That said, with the HP c-Class blade design and easy access for repair, this becomes less of a driver and leads to the second reason for moving to diskless blades, which is ease of provisioning and reprovisioning. When no local data is stored on the blade, it looks like every other similarly configured blade, that is, it becomes stateless. Without on-board data, users can begin to manage their blades as a pool of processing resources rather than as distinct individual servers. This also makes a blade expendable if there is a hardware failure — there is no local data to recover. The administrator can walk up to the front of the chassis, see which blade has failed, pull the failed blade, without affecting the processing on any of the other blades, replace the blade and walk away. If you are running VMware VI3 with high availability (HA), (discussed in more detail later), then when the failure was detected, the processing would have been restarted on another blade. When the failed blade had been replaced, processed would move back onto the new blade automatically. This is the beginning of true utility computing.

Diskless blades get their boot environment over the network and off the SAN using a protocol called Preboot Execution Environment (PXE). When a server blade is booted, it will send out a request, not unlike a DHCP request, for a boot server. A boot server will respond with the file path to a network bootstrap program (NBP) which the blade will download into memory and begin to boot. Depending on the network, this process could take a few minutes, though hopefully, booting is not a process that is performed very often. Using diskless blades has the added advantage that it allows IT to configure, update, and manage all their operating environments on the storage area network (SAN).

There may also be reasons to use local storage, for example, operations where a SAN would not be used, or environments that require a lower storage cost than SAN storage. To support these and other environments where larger amount of direct-attach storage is required, HP has recently created a specialized Storage Blade, the SB40c, which is placed in the blade chassis slot adjacent to the server blade accessing the storage. Each c-Class chassis has been pre-wired with PCI express channels between adjacent slots. To access the storage blade, a PCI express Mezzanine pass-thru card is placed on the server blade to provide connectivity over to the storage blade. Today, each storage blade can hold up to six hot-swap SFF SAS (Serial Attached SCSI) drives for up to 876 gigabytes (GB) of storage in a half-height blade. They also support up to six 60 GB SFF SATA (serial ATA) drives.

In addition to the above options, most blade systems are also attached to a SAN or network attached storage (NAS). HP estimates roughly a 70% SAN attach rate for blade systems.

Blade Networking Options

Blades support a number of networking options for IP-based networking via Ethernet, Fibre Channel (FC)-based networking and high speed low-latency InfiniBand (IB). Every blade comes with at least one 1-gigabit Ethernet (1GbE) network interface card (NIC) standard for management and lights-out operations and some number of expansion slots for additional network capabilities. These networking options are supported primarily through mezzanine (or daughter) cards that attach to the blade. Each blade has support for some number of PCI express expansion slots, e.g., the HP Proliant BL480c blade server supports three expansion slots. The network signals are carried through the midplane to switch modules, pass-thru modules or to HP's Virtual Connect module on the other side of the midplane (see Figure 5 on page 6).

Ethernet

In addition to the standard 1GbE NIC, HP's c-Class blades have the added benefit of having at least 2 integrated multifunction gigabit NICs. Multifunction NICs have additional support for a TCP offload engine (TOE) to unburden the CPU from performing networking tasks, as well as hardware iSCSI (internet SCSI protocol) initiators to support iSCSI SANs and remote direct memory access (RDMA). These

Architectural View of Internal Chassis Interconnect

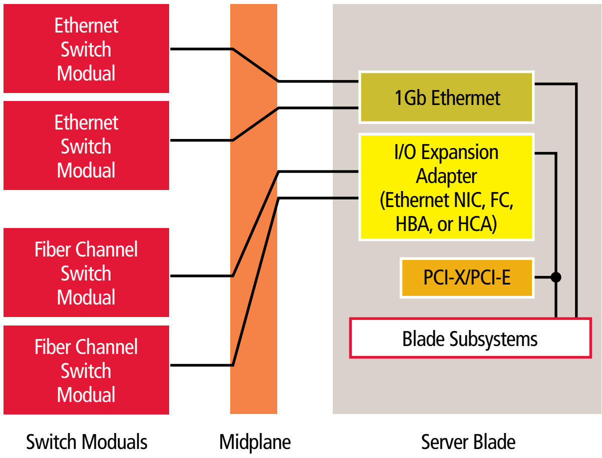

Figure 5

added features all increase the performance of IP network functions. HP also supports both dual and quad-port 10/100/1000 Mbps (megabits per second) mezzanine adapters for applications requiring high network bandwidth.

If Ethernet connectivity and performance was a problem in the past, it is no longer a problem on the newer HP p-Class or c-Class blades. These blades come standard with two 1GbE multifunction NICs and support for additional standard or multifunction 1GbE mezzanine adapters (with 10GbE promised for the future).

On the other side of the midplane, the user must decide what interface module to use. For Ethernet, there is a choice between adding an Ethernet switch module or a Pass-Thru module. The switch module, often the least expensive option when the added cost of cables and additional external switches is factored in, can take inputs from all the blades and consolidate the number of external ports, reducing the number of cables required to access the blades. HP c-Class Ethernet Switch Modules support a variety of options, depending on the module, up to eight RJ-45 uplinks and four optional fibre SX ports with support for layer 2 and one with layer 3 processing. Most blade vendors support several Ethernet switch modules from the top networking vendors, such as Cisco, to provide consistent management across the network infrastructure.

Another option is to use a Pass-Thru module which provides an unmanaged direct connection between the blades and an external network device, such as a switch. This option requires more administrative effort in cabling and may require the purchase of additional external network switches.

Fibre Channel

As with Ethernet, blades support mezzanine Fibre Channel adapters (also known as host bus adapters or HBAs). For resiliency and redundancy, each HBA should have at least two ports. Fibre Channel HBAs can come is 2Gb, 4Gb, and 8Gb bandwidths with the bulk of implementations using 4Gb HBAs. These HBAs are made by the same industry-leaders that provide the PCI bus HBAs.

Either Fibre Channel switch modules or Fibre Channel Pass-Thru modules are placed on the back side of the midplane to move the FC data to external connectors. Like the Ethernet switch module, the FC switch module is the least expensive option and reduces cabling significantly. These switch modules are made by industry leaders such as Brocade and Cisco for compatibility across your storage network. These modules support a minimum of 12 FC ports and can expand up to 24 ports with aggregate switch bandwidth of 192 gigabits per second (Gbps).

TIP: Just as any implementation would have two Ethernet or FC switch modules to provide redundant paths, two Pass-Thru modules should be used for redundant paths.

A Fibre Channel Pass-Thru module avoids having to purchase and configure an embedded FC switch that may be managed and configured differently than your current environment. However, depending on how many spare ports are available on your SAN, there may have to be additional FC switches purchased, along with cables, to connect into your SAN.

InfiniBand

The InfiniBand Switch Module for the HP c-Class enclosure has 16 downlinks to connect to 16 server blades and 8 external uplinks. When using an InfiniBand fabric, a subnet manager is required. For high performance computing (HPC) configurations, HP supports Voltaire Grid Switch family of products. OpenSM from OpenFabric Enterprise Distribution can also be used. HP does not support an InfiniBand Pass-Thru module today.

HP Virtual Connect

A new interconnect option between the server blades and external networks in c-Class BladeSystems (only) is through HP's new virtual I/O offering called Virtual Connect (VIRTUAL CONNECT). Today, there are two Virtual Connect modules that can be ordered from HP, the Virtual Connect Ethernet Module or the Virtual Connect Fibre Channel Module. The Virtual Connect module(s) slide into the back of the c-Class chassis, just like the FC and Ethernet switch or pass-thru modules. The Virtual Connect Modules present MAC and IP addresses, for the Ethernet Module, or world-wide names (WWN) for the FC Module on the external ports. Without Virtual Connect, every time a new blade is added or swapped out of the chassis, both a network administrator and a storage network administrator has to be involved to hook up the blade to the external LAN or SAN. With Virtual Connect, the chassis is configured once, and all further blade configuration changes can be made with no effect on the external LAN or SAN.

This allows the BladeSystem administrator to add new blades, replace failing blades or move blades around within the chassis without having to reconfigure the network external to the c-Class BladeSystem chassis. The Virtual Connect interface to the outside world remains the same allowing the blades to be fungible. All management of connection changes is done through the Virtual Connect Manager, included with every module.

The recently released Virtual Connect Ethernet module supports two 10GbE CX-4 uplinks and eight 1GbE uplink ports externally. Internally, sixteen 1GbE downlinks connect to up to 8 blades though redundant midplane paths. Management features include HTTP, HTTPS, SSL, and support for PXE, WOL (wake on LAN), port VLAN, VLAN Tagging, VLAN pass through, NIC Teaming and port aggregation.

The Virtual Connect FC module supports four external 4Gb FC uplink ports with HBA Aggregation using Network Port ID Virtualization

(NPIV) and sixteen internal 4Gb auto-negotiating FC downlinks connected to up to 4 blades through redundant midplane paths. Management features include support for PXE, port VSAN, VSAN Tagging, VSAN pass through, Link Aggregation Protocol, mirrored profile database and multi-path heartbeat between redundant modules.

Practical Advice

The practical question remains, can four-to-twenty (or more) application environments, together with their operating environments share a server blade without overwhelming the server or the shared I/O components? If the implementation is based on current information and technologies, and well planned then the answer can be a definite yes.

TIP: Planning for server virtualization and blades goes hand-in-hand. When implementation plans call for server virtualization on blades, order blade configurations that will best support the multiple application environments required.

The Plan

As with any implementation plan, the challenge can be in the details. It is, of course, possible to overrun the capabilities of any system, blade or otherwise, by not understanding your applications and loading too many virtual environments with the same or similar usage profiles onto one server. Fortunately, with a little planning this can be avoided.

It is also important to state that not all application environments do well virtualized but may still do well on blades. Specifically, large database applications continue to perform best on their own server. Some customers have moved their large database applications to blades for the technology and management benefits that blades provide but do not use virtualization on that blade. For example, Oracle 10g RAC customers find blades ideal for increased performance as well as space savings.

Before anything else, IT must profile the application environments to determine what the systems requirements are, over the processing life of the application. Some applications have daily, monthly, even yearly spikes in activity. Some spikes are not predictable, caused by some external event. However, it's important to come up with a best estimate for day-to-day operations and deal with anomalies as they arise. Be sure to look at processor, memory, network and storage access patterns.

TIP: Be sure to include operating system overhead as well as application requirements when planning workloads for your virtualized environment.

If the environment being evaluated doesn't already have tools to profile applications environments, many vendors provide consulting services that can help with this process (for example, VMware's Capacity Planner Assessment Service). Once the applications environments have been profiled, set a target for how many environments will run concurrently on the same server. Current users have implemented anywhere from 1 VM per server to 20 VMs per server with an average around 5-6 VMs per server, depending on the environment. With this in mind, choose which applications would be compatible sharing resources. Try not to combine applications which have similar workload profiles as these are the more likely to experience performance degradation due to resource conflicts.

These application groupings become the basis for determining what blade server configurations are required, i.e., how many processors, how much memory, how much network and storage bandwidth is required to meet your needs. Most blade servers today offer many configuration choices, with largely the same processors, memory sizes and I/O options as rack mounted servers. (Note that this was not always the case, fueling many rumors that are still in circulation even though the "problem" is no longer an issue).

When asked what advice they would give to someone just beginning to look at blades and virtualization, the response was, "I'm a big fan of VMware and I have changed my mind on blades. I would like to switch everything to blades and virtualize it all."

It is also important to understand what, if any, constraints are required by the Operating system and/or virtualization software. For example, in VMware VI3, a VMFS can support up to 16 ESX servers (32 if the environment is being replicated for HA). This would mean one BladeSystem c7000 populated with 16 half-height blades, all running ESX servers could use the same VMFS and any virtual machine running on those blades could move to any other blade within the chassis.

Migrating from p-Class to c-Class

For users who implemented the HP p-Class BladeSystem chassis previously, and are now implementing VMware, it is a good time to migrate to the c-Class. The features of the c-Class BladeSystem and blades are more advanced, in a variety of ways, than that of the p-Class, enabling support for improved I/O and more virtual machines per blade. However, this will involve a migration process rather than a simple swap-out upgrade. The advantages of the c-Class infrastructure, which you can expect to leverage over the coming years, are well worth the trouble and HP has services available to ease such a transition. If VMware is already running on the p-Class blades and it is decided to upgrade to c-Class, there will still be a migration process but it will not include the additional migration from physical servers to virtual servers.

Migrating from physical to virtual (P2V)

Don't let the job of migrating your environment from the physical servers to virtual servers be daunting. With a little up-front planning and an understanding of the environment, the transition can move quickly.

First of all, one must decide if the VMware ESX Server VMFS will be created on a SAN or on local disks. Today, over 75% of VMware customers use a SAN for the VMFS. If a SAN will be used, verify that the SAN storage and HBAs are supported by checking the VMware SAN Compatibility Guide http://www.vmware.com/pdf/vi3_san_guide.pdf. To get the most up-to-date compatibility information, go to the VMware website www.vmware.com and search for "SAN compatibility."

When an application is moved from a physical server to a virtual server using VMware's Converter, the application environment including the application data will be moved from the physical server over to the ESX server. That means the data will be copied from the current disk space over to the disk space associated with the VMFS. Because of this, while the application data is being transitioned,

TIP: Unused storage space will be required when transitioning from physical to virtual.

twice the amount of storage will be required. The good news is that once the data has been moved for that application environment, the previous storage can be reclaimed and used again.

It is possible to configure the VMware Converter software to ignore specific LUNs associated with application data during the P2V conversion. This is called Raw Device Mapping (RDM). It may not make sense to migrate the data to the VMFS if the application environment is associated with a very large amount of data, or for performance reasons it may be decided to maintain the data outside of the VMFS. Once the P2V is finished, the LUN can be reconnected to the VM by creating an RDM. However, the VM loses its mobility. This also makes for more complicated management of the ESX server. If RDM is required for an application, consider whether that application is better left on a physical system.

Keeping up with the times

Five years ago, blade servers and server virtualization were just emerging for the innovators and early adopters who liked the bleeding edge. Today, both virtualization and blade technologies are becoming main-stream, implemented successfully from Fortune 100 enterprise organizations to forward thinking SMBs. The markets for both technologies have seen significant change over this time. Change in capabilities, change in leadership among the top vendors and change in how users implement and manage these technologies within their infrastructure. It also means added capabilities for increased performance, bandwidth, and ease of management.

TIP: The rate of change in these technologies has been significant over the past 5 years. Avoid bad business decisions, by making sure your implementation staff and partners are current on the latest tools, technologies and trends on both blades and virtualization solutions.

One example of this can be seen over the lifetime of the HP BladeSystem and blades. Initial blades had limited network bandwidth support. As such, these blades would not have made good candidates for multiple virtual environments, especially environments with high levels of network access. As customers required more bandwidth, HP developed new p-Class blades with greater bandwidth. Additionally, even more bandwidth is supported in the latest HP c-Class blades and chassis. Many new advances have been made in I/O bandwidth, power and cooling, and now in virtual I/O through Virtual Connect.

When the first release of VMware ESX became available, neither Intel nor AMD had virtual-assist technologies to help with a hypervisor's privileged status like they do today. With every release of ESX, VMware works to improve overhead of the virtual environment, especially with regard to I/O. Before implementing VMware, it is imperative that the compatibility of both the system and the storage is verified by checking the VMware website.

Cost Justification

As part of the strategic decision to move to blades and virtualization, the CIO and probably the CFO will want to see financial justification. To assist in this process, CDW has tools, including an ROI analysis tool and a TCO analysis tool. The outputs from these tools help show the bottom-line value of moving to blades and virtualization.

Today, IT is required to provide cost-effective service to an environment that is made up of hundreds of servers consuming power, creating heat, requiring cooling, cabling, security, server, storage, data, and software management, with servers running less than 10% utilization and generally running only one application.

Virtualization and blades offer a reliable, flexible and substantially more cost effective alternative. With HP and VMware solutions, users are consolidating up to 100 servers into a single rack with up to sixteen blade servers sharing power, cooling, IP network, SAN network cabling, disk storage, and management, reaching server utilization rates at 80% or above. The combination of VMware's Virtual Infrastructure 3 (VI3) and HP's c-Class BladeSystem offer significant value for IT today and into the future.

About the Authors

Anne Skamarock, advisory analyst with Focus Consulting, has spent nearly 30 years in high-tech fields in various positions, including end user, systems administrator, scientific programmer, NFS software engineer, backup and recovery software engineer, technical sales, marketing, and product management, and industry analyst with companies including SRI International, Sun Microsystems, Solbourne Computer, StorageTek, and Enterprise Management Associates (EMA). She is co-author of *"Blades Servers and Virtualization: Transforming Enterprise Computing While Cutting Costs.* She also co-authored Storage Solutions: A Buyer's Guide and has written regular columns for NetworkWorld and TechTarget, and numerous white papers.

Barb Goldworm, president of Focus Consulting, has spent thirty years in the computer industry, in various technical, marketing, sales, senior management, and industry analyst positions with companies including, IBM, Novell, StorageTek, Enterprise Management Associates, and multiple successful startups. A frequent speaker at industry events, Barb currently chairs the *Server Blade Summit* conference and is a regular speaker and expert columnist for *TechTarget SearchServerVirtualization*, having also written regularly for ComputerWorld SNWOnline, and NetworkWorld. She also co-authored *"Blades Servers and Virtualization: Transforming Enterprise Computing While Cutting Costs."*

Focus Consulting (www.focusonsystems.com) is a research, analyst, and consulting firm focused on systems, software and storage. Focus offers a team of industry veterans combining decades of experience in systems, software, storage, management and services for both the enterprise and SMB markets. Focus analysts bring expertise from technical, marketing, business, and executive roles; have extensive experience in market research and analysis, have written regularly for leading industry publications; have authored and edited books; and are frequent speakers and educators at industry events. The Focus team brings a comprehensive yet focused view of how to meet demanding business needs in this world of ever changing technology, while keeping an eye on ROI, TCO and bottom line value.

NEW

Enclosure shown with both 460c and 480c blade servers; all sold separately

HP ProLiant BL460c Blade Server

Call your CDW account manager for complete Blade configuration

- The dual-core Intel® Xeon® Processors 5140 (2.33GHz) allows a greater workload per processor.
- Memory-intensive applications benefit from the 2GB std., 32GB max. (PC2-5300) memory buffers allow faster speeds and more capacity.
- With up to two SATA/SAS hard drive bays, you get quicker access to data and increased storage bandwidth.
- On-chip cache (4MB Level 2 Cache) increases dual-core performance and provides faster processor hit rate.

Call for pricing CDW 1007066

HP Care Pack Three-year, four-hour, 24 x 7, onsite **Call for pricing** CDW 1005222

intel Xeon® inside™ Dual-core. Do more.

HP StorageWorks D2D Backup System

The HP StorageWorks D2D Backup System provides reliable, consolidated data protection for up to four servers in a single, self-managing device, requiring less manual handling (benefiting organizations with limited IT resources). It increases the reliability of your backups by reducing the risk of human error in managing tape drive and media hardware.

This disk-to-disk backup solution lets you easily and quickly restore lost or corrupted files from online backups. The HP StorageWorks D2D Backup System works with your backup software application to provide daily automated backup for multiple servers at less than the cost of a single LTO tape autoloader. It also supports the direct backup of up to four servers simultaneously.

- 1Gb iSCSI (Ethernet), optimal (also supports 100 base-T)
- Emulates HP LTO Ultrium 2 Tape Drive, HP 1/8 LTO Ultrium 2 Autoloader
- Maximum number of virtual tape cartridges emulated 96 (24 per tape device, up 200GB per cartridge)

HP StorageWorks D2D120 2TB D2D **Call for pricing** CDW 1136418

HP StorageWorks D2D120 2TB D2Dwith DPX **Call for pricing** CDW 1136419

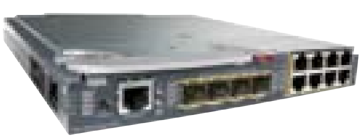

HP Cisco® Catalyst® Blade Switch 3020

Flexible to fit your needs, the HP Cisco® Catalyst® 3020 Blade Switch for HP c-Class BladeSystem provides HP customers with an integrated switching solution from Cisco Systems®. This solution provides Cisco resiliency, advanced security and enhanced manageability to the server edge while dramatically reducing cabling requirements.

- Provides wire-speed switching on 16 internal ports
- Features CiscoWorks Software which includes multilayer configurations such as routing protocols, Access Control Lists (ACLs) and Quality of Service (QoS) parameters
- Supports Cisco's Identity-Based Networking Services (IBNS) which prevents unauthorized users access into your network

Call for pricing CDW 1000407

VMware® Infrastructure Data Center Management and Optimization Suite

VMware Infrastructure is a software suite for optimizing and managing industry-standard IT environments through virtualization — from the desktop to the data center. VMware Infrastructure provides built-in management, resource optimization, application availability and operational automation, delivering transformative cost savings and increased operational efficiency, flexibility and service levels.

Call for pricing

To learn more or discuss your virtualization needs, visit us at www.cdw.com/solutions or call your Dedicated Account Manager today at 800.800.4239.

033007 34377