

WHITE PAPER

Improving Economics of Blades using VMware



Improving Economics of Blades with VMware

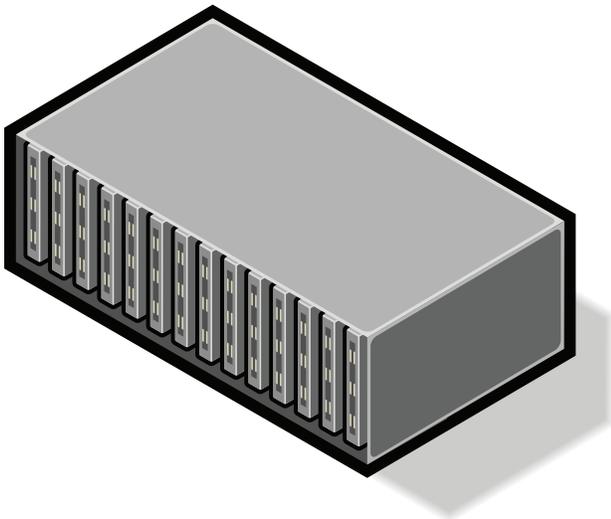
Executive Summary

Today IT efficiency is critical for competitive viability. However, IT organizations face many challenges, including, growing capacity while conserving or reducing operations cost, improving service availability and change request response time and increasing consistency and predictability.

VMware® ESX Server™ with VirtualCenter™ deployed on blade servers is an effective solution to these challenges. The combination of VMware unique manageability benefits, maximum possible density, server, storage, and networking virtualization technologies enables IT infrastructure that is highly optimized, cost-effective, and consistent.

VMware software further enhances density, operational efficiency, and easy provisioning benefits of blade servers. In addition, VMware software provides more flexibility and responsiveness in provisioning new software services and maintaining existing services. VMware VirtualCenter deployed on blade servers manages virtual and physical infrastructure to let you continuously consolidate workloads for optimal hardware utilization and decreased operational costs. With VMware you can:

- Save 35% or more in operational costs due to higher systems utilization
- Make your infrastructure utility-computing ready with on-demand dynamic provisioning, affordable availability, and flexible management.



Introduction

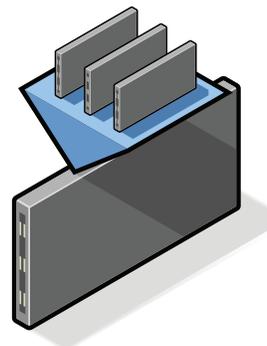
IT managers are constantly under pressure to increase operational efficiency. Ever increasing requirements for maintaining more applications, supporting a larger user base and providing better uptime do not come with increased IT budgets, higher headcount, or larger facilities.

To contain costs while growing their data centers, IT managers turn to consolidating computing resources into a smaller number of physical and logical elements and automating allocation, deployment, and management of these elements.

Consolidation strategies include consolidating heterogeneous servers to a central physical location, scaling up to larger servers, scaling out to identical servers, and native application consolidation. Each approach has its own “sweet spots” depending on the required application mix and constraints. For example, an 8-way server running a VMware virtualization platform is well suited for an environment where many varied applications need to be hosted and some of the applications are legacy or ill behaved. Native application consolidation may be a better approach for large uniform loads such as large-scale databases.

In this paper we will focus on consolidation with blade servers. Such consolidation yields very good results when IT is looking to improve server manageability and optimize operational parameters such as provisioning, cabling complexity, power and cooling expenses, and floor space utilization, while containing costs, and increasing system uptime.

VMware virtualization and VMotion™ technology provide unique provisioning and management benefits: complete virtual machines can be automatically deployed using centrally stored template libraries, virtual machines can be migrated between physical blades with zero user downtime, computing resource allocation can be dynamically managed without interrupting the workload. These benefits further enhance benefits of consolidation solutions on blades servers.



Blade Server's Benefits and Limitations

New blade server offerings such as IBM BladeCenter™ and HP ProLiant™ BL p-class servers go a long way to address the consolidation needs of enterprise IT managers.

A blade server typically includes a rack-mountable chassis of 6U to 8U size that can be populated with eight to fourteen blades. Each blade is an independent server with one or two processors, its own RAM, local storage and NICs. Usually, a blade is used to support a single application workload at a time. The chassis incorporates elements shared by all blade servers such as built in network switches, redundant hot swap power supplies, fans, I/O and management modules. The form factor, high density, and standardization in blades environment are conducive to more robust system architecture with consolidated I/O and shared external storage such as SAN.

Let us consider two 42U racks deployed side by side in a data-center. The first rack is populated with six blade server chassis and yields very high processor density, as high as 168 processors per rack. The second rack is populated with 42 1U rack optimized two processor servers yielding only 84 processor per rack density.

Such high density is invaluable if the datacenter floor-space is limited but there is a need to add computing power. In addition, blade server solution would provide 7% to 12%¹ savings in power and cooling costs compared to non-blade solution.

Deploying blade servers is a good approach when the IT organization is constrained by people resources. Managing six distinct physical elements rather than forty two provides for better platform stability and is less costly. In addition, blade servers are optimized for easy cable management both via integration with cable management systems and through utilization of built in elements such as switches and I/O modules.

New server provisioning is also simplified. Often the IT administrator uses a "scale out exactly" approach where an existing blade is cloned and plugged into an available slot on an existing chassis.

Deployments based on blade servers also provide better fault resiliency because many fault prone elements such as power supplies and I/O modules are now redundant. Such redundancy is cost prohibitive in the non-blade solution because such service modules are dedicated rather than shared.

However, despite all the benefits of blades, IT managers should keep in mind certain limitations:

- Blade server cost efficiency declines sharply if its density potential is underutilized, for example, when chassis are not fully populated or single processor blades are used.
- Local storage availability on blades is limited.
- Few expansion slots are available for either HBA adapters or other peripherals such as additional NICs or encryption accelerators. Deploying peripherals may reduce blades density.
- Cost of deploying blades is a step function. The cost jump is especially painful when deployment of a new chassis is required.

VMware ESX server with VirtualCenter enhances many of the blades benefits and mitigates many of blades limitations.

"We are also experiencing huge cost savings from the implementation of VMware, a software application that consolidates server infrastructure. This application allows us to run two applications on one server that, without the software, would each require its own server. As we implement it in each data center, we have spare servers that can be deployed for new applications. Besides the cost savings on capital outlays, having fewer servers provides cost savings in headcount and space"

*Michael A. Nelson, Finance Manager,
Corporate Information Technology, Pfizer, Inc.*

VMware Improves System Utilization and Operational Flexibility

For the last five years VMware virtualization technology has enabled IT customers to reduce the costs of their IT operations, providing economic, operational and business flexibility benefits. VMware ESX Server transforms a physical system into a pool of logical computing resources. Operating systems and applications are isolated in secure and portable virtual machines. System resources are dynamically allocated to any operating system based on need, providing mainframe-class capacity utilization and control of server resources. VMware technology enables fine grain resource management².

Advanced resource management control allows IT administrators to guarantee service level metrics across the enterprise.

ESX Server is widely deployed throughout IT organizations of Fortune 500 enterprises. With VMware ESX Server:

- Applications running on dedicated systems can be moved into separate, completely isolated virtual machines on a single system of higher density.
- Servers can be remotely managed from any location, simplifying server maintenance and contingency planning.
- Service levels can be guaranteed with advanced resource management controls.
- Common monitoring and management tasks are scriptable.
- Capacity can be increased without adding new physical systems.

To improve hardware utilization IT managers can take advantage of the fact that workloads are sometimes idle and that different applications are bound by different hardware resources (i.e. some applications are memory bound, some are CPU bound) and that peak usage does not occur at the same time with different workloads.

VMware VirtualCenter product introduced in November 2003 enables efficient application of virtualization technology to scale-out solutions as well as scale-up solutions. VMware VirtualCenter gives IT manager the ability to view server farms as a unified pool of CPU, memory, and storage resources and to allocate these resources on demand. With the Centralized Management Console, administrators can monitor computer system availability and performance as well as manage all their virtual machines. VMotion transfers the entire system and memory state of a running virtual machine from one ESX Server to another without interrupting user sessions for continuous workload consolidation and zero down time 9 to 5 maintenance. VirtualCenter also integrates with market leading platform management software to provide IT managers with the greatest possible control over datacenter infrastructure.

Deploying ESX Servers with VirtualCenter and VMotion in conjunction with blade servers provides all the benefits described above on a robust high-density datacenter-optimized platform.

“With VMware we gained efficiencies on many levels, from optimizing the utilization of hardware resources and consolidating server platforms, to establishing more effective platforms for testing and business continuity / disaster recovery”

*Warren Jones, VP, IT Operations,
Guardian Life of America*

VirtualCenter and VMotion extend Virtualization Benefits to Scale-out Hardware Farms

As with any server, blade servers come in different configuration flavors that affect both computing power and acquisition cost. For example, a low-spec configuration may be a single processor server with 512MB of memory and a single IDE drive which lists for \$3000. A high-spec blade may be configured with 2 Intel Xeon 3GHz processors, 2GB of memory, a SCSI drive, and ability to connect to external SAN and be priced at \$6300³.

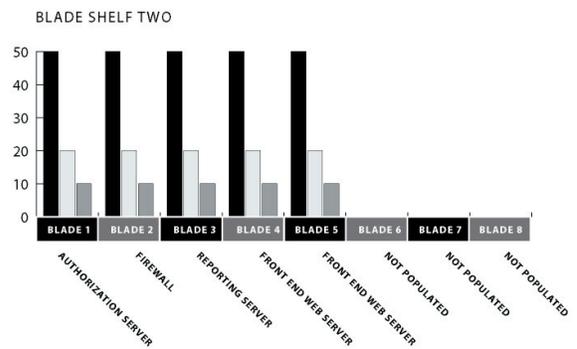
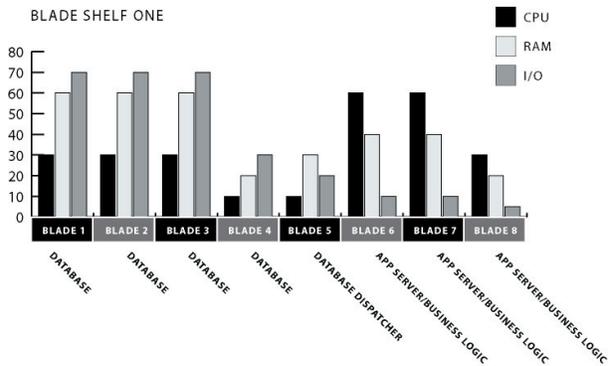
Virtualization technology enables high hardware utilization and allows the solution to enjoy the robustness of high-spec configuration without incurring the cost penalty.

Let us consider blade server implementation of a 3-tier web application requiring a database, application platform middle-ware with business logic, redundant front-end web servers and supplemental servers for authorization, firewall, and reporting.

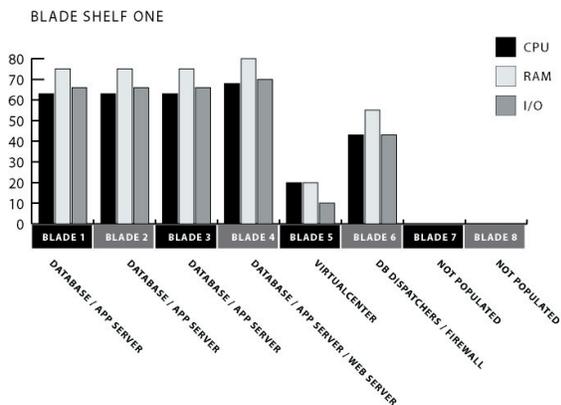
In the first example the complete 3-tier application deployed on blade servers requires seven high-spec blades and six low-spec blades taking up two blade shelves.

In the second example, VMware ESX Server is deployed on top of blades. In this example the complete 3-tier application takes up only five high-spec blades. One additional low-spec blade may be used for VirtualCenter, if VirtualCenter is not yet deployed in the datacenter. The more consolidated solution with VMware takes advantage of the fact that database and application servers are bound by different computing resources: database is memory and I/O bound while application server is CPU bound.

BLADE SERVERS WITHOUT ESX SERVER & VMWARE VIRTUAL CENTER



BLADE SERVERS WITH ESX SERVER & VMWARE VIRTUAL CENTER



Due to high storage usage by the database, such a solution needs to be deployed in conjunction with SAN. However, low utilization of the blades in blade chassis two in the deployment without VMware software makes it cost prohibitive to populate that chassis with blades that are SAN enabled. As a result, part of the network does not enjoy the robustness that SAN provides. Further, deploying applications on local storage has a significant physical element that makes administration more expensive. To exacerbate the situation, two different business processes are required for maintenance, back-up and contingency planning of this implementation.

In the second example all virtual machines and virtual machine templates reside on SAN. This practice encourages standardization, promotes best practices, and greatly simplifies contingency planning and disaster recovery.

VirtualCenter with VMotion enables instant deployment and computing resources re-allocation, provided unused computing resources are available. It is even possible to deploy unattended automatic script-based workload migration to meet performance and availability metrics.

When it is necessary to deploy an additional blade, such deployment is also simplified because all blades have uniform hardware configuration and are installed with ESX Server software, while in the first example blade hardware and software configuration varies depending on the specific load planned for this blade. Deployment with VMware has better consistency and simplifies capacity planning. In addition, uniform blade configuration means that fewer spare hot-plug blades need to be in stock to maintain the same level of fault resilience.

	1U rack servers	Blades server	Blades with VMware
Utilization	5%-15%	5%-15%	60%-80%
Operational costs		25% reduction	85% reduction
Downtime for service and provisioning	4 hours	1 hour ⁴	Zero downtime
Work required to adjust workloads	Requires shut-down, re-cabling, and loading new workload	Requires shut-down and loading new workload	Continuous, zero downtime

VirtualCenter enables efficient centralized management of both physical and virtual computing resources.

VMware ESX Server with VirtualCenter deployed on blade server has reduced the number of blades required from 13 to 6, thereby reducing number of required chassis from two to one.

Such reduction provides very significant saving in operational cost:

- 50% facility space saving.
- Close to 50% reduction in power and cooling expenses.
- 50% or more savings in physical administration costs such as re-cabling, inventory, etc.

Blade Migration and Managing Heterogeneous Environments

VMware technology also simplifies migration to blades. Using VMware P2V utility, an IT manager can capture an existing physical machine complete with operating systems, patch level, anti-virus utilities and the workload application and migrate it exactly to a virtual machine. What's more, an administrator would then select one configuration per workload and create a "gold" template that could be replicated exactly to scale out. As the processing needs grow, additional copies of this virtual machine can be provisioned remotely within minutes.

Rather than moving all the workloads over to blades servers at once, VMware allows gradual migration.

Hybrid environment of both blades and non-blade servers can be managed by VirtualCenter as unified pool of computing resources. Virtual machines can be moved between blade and non-blade servers with VMotion without interrupting user sessions.

Blade Virtualization Compared to Blade Server Physical Consolidation

On the surface it appears that low-spec blades are so inexpensive that it makes sense to just use low-spec blades and deploy one workload per blade. Let us compare the three year cost of two deployments: one where low spec blades are used and the other where high-spec blades are used with VMware software.

Firstly, let us note that the acquisition costs of low-spec blades for SAN environment are about \$4000 which includes the cost of blade itself and chassis cost allocation. Considering each blade runs a single workload, this is cost prohibitive. Thus, the low-spec blade solution requires a tradeoff of either low cost or high robustness.

High-spec blades have on average two times compute power vs. low spec blades (assuming two 3GHz Xeon processors with 2Gb memory vs. one 2.6 GHz Pentium IV processor with 512Mb of memory). Assuming a single workload utilizes 30% of low-spec blade computing power and allowing for the possibility that 50% of the workloads can peak simultaneously doubling their usual computing resources utilization, we can consolidate five workloads onto a single high spec blade with VMware ESX Server for an average utilization of 67%. Let us note that this is a conservative ratio. The target utilization rate can be higher with VMotion technology which enables automatic continuous workload consolidation and would be able to relocate a workload with unexpected resource requirements to a different available hardware resource.

This example clearly demonstrates that VMware software provides substantial economic benefit especially in operational costs of power, cooling, and floor space. In addition, VMware software is based on more robust SAN technology and provides more flexibility to system administrator.

	HS blades w/VMware	LS blades
Number of workloads	3000	3000
Number of blade chassis	75	375
Number of racks	13	63
Estimated hardware acquisition cost	\$6.5MM	\$12MM
VMware license cost	\$1.8MM	N/A
3-year cost of floor space, power, cooling	\$1.7MM	\$8.5MM
3-year hardware and software support and maintenance	\$16.7MM ⁵	\$29.3MM
3-year cost to run workloads	\$24.9MM	\$49.9MM

Conclusion

VMware ESX Server with VirtualCenter deployed on blade server builds on the core benefits of blade servers to further improve computing density and make IT infrastructure more robust. Virtualizing blade server resources promotes use of blades with the highest possible density, and, thus, further reduces datacenter operational costs by as much as 70% compared to operating the same workloads on non-virtualized blades. VMware VMotion technology enables greater IT flexibility including faster response time, no-touch provisioning, and zero downtime maintenance. VMware software makes SAN deployment more affordable. Deploying VMware software promotes standardization, easier contingency planning, and best business practices.

IT managers who plan to deploy blade servers should definitely consider deploying VMware ESX Server with them. The VMware Sales Team can help your IT organization determine how VMware software will provide these benefits in your particular datacenter environment. Using ROI tools, case studies, and other tools, VMware will work with you to design and implement specific success criteria so you can evaluate our software effectively. Visit us on the Web at www.vmware.com, email us at sales@vmware.com, or call us at 877-4VMWARE to get started.

(Footnotes)

1 IBM whitepaper "To blade or not to blade... that is the IT question" by Mark T. Chapman and Don Roy, Sept 2003

2 Fine grain resource management means that a fraction of a resource can be assigned to a particular workload. For example a workload may be assigned 20% of one CPU and 1.5 (teamed) NICs.

3 For the purposes of this discussion we use list prices of hardware. Actual acquisition price may include volume or channel discounts.

4 We assume blade ghosting /cloning is employed

5 The support cost estimate includes fees for support of VMware software

V00014-20001205



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
Copyright © 2004 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242 and 6,496,847;
patents pending. VMware, the VMware "boxes" logo, GSX Server and ESX Server are trademarks of VMware, Inc. Microsoft,
Windows and Windows NT are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus
Torvalds. All other marks and names mentioned herein may be trademarks of their respective companies.

