White Paper

Technical and Commercial Comparison of Citrix XenServer and VMware vSphere

November 2011

Customers exploring server virtualization solutions should consider the technical and commercial aspects of Citrix XenServer and compare them to VMware vSphere

Overview

This document provides a technical and commercial comparison of Citrix® XenServer® and VMware® vSphere™, two of the leading server virtualization products on the market. This comparison illustrates that XenServer is a powerful and feature rich alternative to vSphere and distinguishes itself in a number of key areas including open architecture, performance, storage integration and total cost of ownership (TCO).

- **Citrix XenServer 6** is a free virtualization platform based on the open-source Xen® hypervisor. It includes Citrix® XenCenter®, a multi-server management console providing management of critical operations such as virtual machine templates and snapshots, shared storage support, virtual networking, resource pools and Citrix® XenMotion™ live migration. Citrix offers advanced management capabilities in the premium editions of XenServer: Advanced, Enterprise and Platinum.
- **VMware vSphere 5** is a server virtualization platform that includes only the VMware® ESXi™ hypervisor and associated management tools. vSphere is available in three editions: Standard, Enterprise and Enterprise Plus. vSphere servers are managed from vCenter Server, VMware's multi-server management console. VMware requires that management via vCenter Server be purchased separately. Advanced management features such as vCenter Server Heartbeat (high availability for vCenter) and Site Recovery Manager also require independent licenses. VMware does not support its ESX hypervisor for vSphere 5 deployments, and has published several migration guides to assist customers with their migration to ESXi.

Comparisons

System Architecture

Both products feature bare-metal or Type 1 hypervisor technologies that are installed directly onto physical servers without requiring a host operating system (OS). It is widely accepted that this bare-metal approach offers significantly better performance and manageability than solutions reliant on a host OS, such as VMware Server® (GSX).

VMware ESX Architecture

VMware ESX was a first-generation architecture that predated virtualization-aware operating systems, including most Linux® distributions and Microsoft® Windows Server® 2008, and processors such as Intel® VT and AMD-V™. Versions of vSphere prior to vSphere 5 included an option to leverage the ESX hypervisor. vSphere 5 now only supports the ESXi hypervisor.

VMware ESXi Architecture

VMware created the ESXi hypervisor as its first second generation solution. ESXi is designed to leverage the virtualization features which emerged in Intel and AMD processors starting in 2005. These virtualization features were designed to address several of the limitations in the Intel x86 architecture, and by extension eliminate the need for solutions like binary translation. Since VMware had invested many years tuning its binary translation algorithms, it was not uncommon to find that they occasionally out-performed the initial implementations of the Intel-VT and AMD-V replacements. As a result, VMware marketed ESXi as a free offering, and initially it had a limited HCL.

Architecturally, ESXi also differs from ESX in management functions. An ESX based deployment will often leverage the ESX Service Console and management agents are installed in this console. In contrast, ESXi uses an API based management model, and lacks a service console. This difference in management model makes transitions from ESX to ESXi challenging. In 2010, VMware announced that the last vSphere version supporting the ESX hypervisor will be vSphere 4.1. vSphere 5 fulfills that promise and as a result, customers running ESX will be faced with a potentially destabilizing event should they decide to upgrade to a future version of vSphere. To mitigate this, VMware has published several documents detailing migration best practices.

Citrix Architecture

The architecture of XenServer is quite different from vSphere, as XenServer development coincided with the availability of virtualization-aware Intel and AMD processors and operating systems. XenServer is built on the open-source Xen hypervisor, which is also the basis for nearly two dozen commercial virtualization products and the engine powering the world's largest virtualization deployment, the Amazon™ Elastic Compute Cloud™. Xen itself is fully supported in mainline Linux as both a guest (2.6.27+) and as a control domain (2.6.39+) meaning that Linux fully supports Xen.

Rather than using binary translation like VMware, XenServer uses a combination of paravirtualization and hardware-assisted virtualization. XenServer was the first solution on the market to employ paravirtualization, which allows a guest operating system to be fully aware that it is being run on virtualized hardware. This collaboration between the OS and the virtualization platform enables the development of a simpler, leaner hypervisor that offers best-in-class performance.

Today, XenServer supports paravirtualization with a number of Linux distributions, including Red Hat® Enterprise Linux®, Novell® SUSE, Debian®, Oracle® Enterprise Linux and CentOS. For guest operating systems that can't be fully paravirtualized, such as Microsoft Windows®, XenServer is designed to leverage hardware virtualization assist technologies, widely available on today's Intel VT and AMD-V processors.

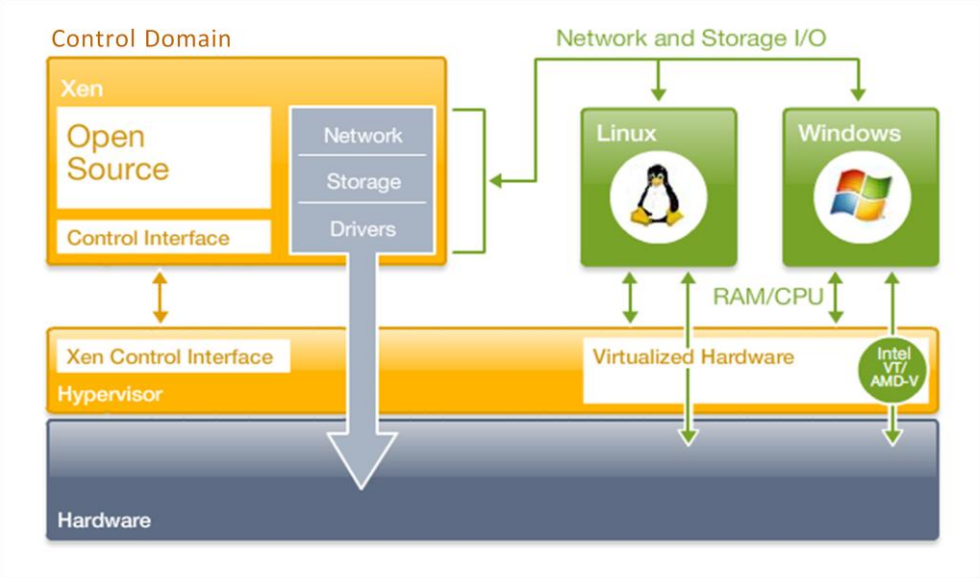

Figure 1 - XenServer Hypervisor Architecture

Despite initially downplaying the merits of paravirtualization, VMware adopted it within ESXi in the form of its VMI technology. The release of a VMI performance white paper further highlights VMware's belief that paravirtualization is superior to binary translation. Unfortunately, thus far, only a few Linux operating systems (specific versions of SUSE and the community-supported Fedora) are enabled to run with VMware VMI on vSphere.

Device Drivers

The XenServer approach to device drivers is also significantly different from the one used by VMware. With XenServer, all VM interactions with the hardware are managed through the *Domain 0* control domain, which itself is a specially privileged VM running on top of the hypervisor. The XenServer Domain 0 is shown in the upper-left-hand side of Figure 1 and runs a hardened, optimized instance of Linux. It's important to recognize that to the administrator, Domain 0 is part of the overall XenServer system and requires no additional installation or management. Domain 0 enables XenServer to leverage standard open-source Linux device drivers, resulting in extremely broad hardware support. Because of this design, XenServer can even run on laptops or workstations—systems on which VMware's proprietary drivers are much less likely to function.

Installation, Configuration, and Administration

XenServer has a straightforward installation process, often referred to as the *10 minutes to Xen* experience. XenServer is installed on the host systems using a CD or network-based installation process. The XenCenter GUI-based administration console is then installed on any Windows® based computer. System configuration information is kept in an internal data store within the XenServer control domain and is automatically replicated across all servers that are managed together (forming a *resource pool*). The resource pool configuration provides highly available core management services that run on a management architecture with no single point of failure. This architecture avoids the need for a separate database server for the core management functions.

Similar to XenServer, the VMware ESXi hypervisor is installed on the host servers. For management and configuration VMware uses vCenter Server, which runs as a Windows service on a separate management server. Unlike XenCenter, vCenter requires a third-party database for storage and management of host system configurations. For redundancy and availability of the core management services, VMware recommends clustering software such as its vCenter Server Heartbeat add-on product.

Guest Operating System Support

VMware and XenServer both support all versions of Microsoft Windows which are supported by Microsoft and various Linux operating systems including Red Hat, Oracle, SUSE and others. VMware and XenServer are certified for Windows according to the requirements of the Microsoft Server Virtualization Validation Program (SVVP).

Storage Integration

Storage is one of the most important considerations for server virtualization deployments. Both solutions offer support for storage of VMs on local disks, iSCSI or Fiber Channel-based storage area networks (SANs), or network attached storage (NAS). Both platforms require use of a SAN or NAS to support advanced features such as live migration and high availability.

When architecting a virtual infrastructure, VMware recommends its proprietary VMFS file system for storage (although raw disk options are available). Regardless of the storage type, VMFS is the default storage system for virtual machine disk images. VMFS is a clustered file system that, when used with SANs, allows VMware vStorage to take control of certain storage functions, including provisioning and snapshotting, regardless of which vendor's array is being used.

XenServer takes an entirely different approach to storage. XenServer does not impose its own file system on storage systems, but rather leverages the native storage capabilities more directly. For example, with a file-based shared storage system such as NFS, XenServer VMs are stored directly using Microsoft Virtual Hard Disk (VHD) format. With block-based storage such as iSCSI or Fiber Channel SANs, XenServer extends VHD with the Logical Volume Manager (LVM) standard for volume management.

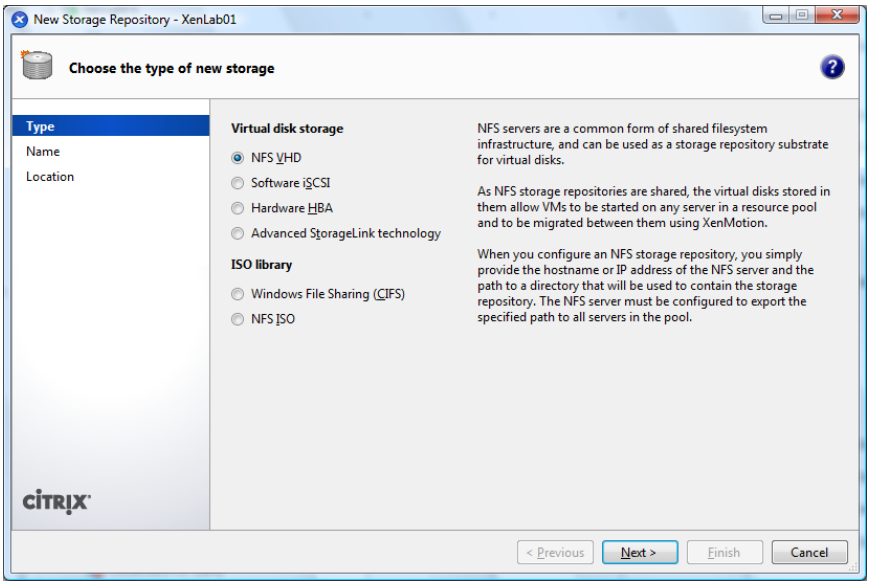

Figure 2 - XenServer Storage Repository Wizard

This differentiated approach to storage integration with XenServer comes in the form of Citrix® StorageLink™. StorageLink enables the virtualization administrator to directly use features of industry-leading arrays such as those from NetApp and Dell EqualLogic. This approach ensures that common storage management process and practices can be leveraged across both virtual and physical environments, and effectively shields virtualization administrators from storage implementation details.

For example, a VM snapshot request within XenCenter will offload this task to the SAN through an interface with the SAN vendor's API. In this manner, StorageLink can optimize performance and permit virtualization administrators to use the advanced features of the array exactly as they were intended, including snapshots, clones and thin provisioning. In addition to vendor-specific APIs, StorageLink supports the ANSI standard for storage management (Storage Management Initiative – Specification), enabling easy integration of many storage vendors' equipment.

Provisioning Services

Citrix Provisioning Services streaming technology allows server workloads (operating system images, applications and configurations) to be provisioned and re-provisioned in real-time from a single shared-disk image. In doing so, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is done on the master image. Provisioning services functionality is ideal for server farms such as Citrix© XenApp™, web servers, application servers and Citrix XenDesktop virtual desktops.

Dynamic workload streaming can be used to rapidly deploy server workloads to the most appropriate server resources (physical or virtual) at any time during the week, month, quarter or year. It is particularly useful for applications that may be regularly migrated between testing and production environments or for systems that might require physical deployments for peak user activity during the business cycle. Provisioning services is a unique feature that only Citrix offers as a feature of its products.

Virtual Machine Protection and Recovery

XenServer and vSphere offer a range of VM snapshot and automation capabilities for backup and recovery needs, including support for traditional agent-based backup solutions from a variety of providers. XenServer improves upon this agent-based approach for Windows VMs through the enablement of a product-specific Volume Shadow-copy Service (VSS) provider. For example, when a backup agent makes a call to the Windows VSS provider, the XenServer VSS provider intercepts the call and redirects it to perform a disk-level snapshot. This disk-level snapshot is a differential copy and as such takes as little as five seconds to complete. In this manner, regular backups can be performed without impact to the performance of the XenServer host and with minimal backup impact of the guest. The VM protection and recovery feature in the Advanced Edition of XenServer builds on this capability by allowing administrators to create policies to automatically take regularly scheduled VM snapshots and then archive the images to a specified storage location.

VMware offers a feature called Data Recovery in vSphere for backup of VMs. Data Recovery is a plug-in for vCenter Server that schedules regular disk snapshots for VMs. Some customers see this as an alternative to traditional agent-based approaches; however, Data Recovery does not offer application-level awareness, which is often cited by customers as a reason to use agent-based solutions for workloads such as SQL Server®, Oracle Database and Application Suite and Microsoft® Exchange® and Active Directory®.

High Availability and Fault Tolerance

During server consolidation efforts, one of the key concerns often is the impact failure of a virtualization host might have to data center operations. High availability ensures critical VMs are automatically restarted on another physical host should the original host running the VM unexpectedly fail. This can reduce the amount of downtime for the workload, as well as eliminate the need for administrative intervention. Both vSphere and XenServer have high-availability features that offer granular policies governing the behavior of specific VMs after a host failure.

Additionally, both Citrix and VMware offer options for fault tolerance, a feature that maintains mirrored instances of VMs running on separate hosts. In the event of a host failure, the mirrored instance can maintain continuity of the workload. With vSphere, fault tolerance is offered in the Enterprise Edition and above and is only supported for VMs with 1 vCPU. Citrix has partnered with several third-party vendors, such as Marathon Technologies and Stratus, to create more robust fault tolerance options for XenServer.

Workload Placement and Balancing

Determining the optimal placement of a VM becomes an increasing challenge as the number of VMs increases. Often load factors on a host or within a VM can create situations where a given resource pool is underperforming relative to its potential. The task of managing the performance of the virtual infrastructure includes not only where best to start a given VM, but also how to optimize the entire pool for normal operations and critical events such as a host failure. These operations are further complicated when sustainability initiatives such consolidating VMs as part of a data center power management plan occur. XenServer offers workload balancing (WLB), which leverages utilization data from the CPU, memory, disk I/O, and network I/O on the hosts and VMs to guide the initial and ongoing host location for VMs. There are two optimization modes for WLB: *optimize for performance* and *optimize for density*. Optimizing for performance ensures that minimum performance thresholds are maintained, whereas optimizing for density places VMs on the minimum number of hosts to reduce power consumption.

vSphere Distributed Resource Scheduler (DRS) is a feature that guides the initial VM placement and partially or fully automates load management of VMs. Unlike WLB, DRS does not allow much customization of the load management algorithm and is based only on CPU and memory utilization. In vSphere 5, VMware introduced Storage DRS, but unlike WLB which monitors the I/O performance for a VM and migrates the VM to a host with less I/O contention, Storage DRS monitors the storage I/O of the virtual disk and migrates the underlying virtual disk to an alternate LUN with less traffic. At launch, Storage DRS was not compatible with vSphere Site Recovery Manager.

Disaster Recovery

Disaster recovery (DR) involves the duplication of virtual server infrastructure and data at remote facilities for recovery in case an event makes the primary site inoperable or inaccessible. Virtualization simplifies disaster recovery in many ways, as server workloads packaged as VMs are easier to transport and restart on remote systems.

XenServer and vSphere support multi-site deployments, whereby VMs can be made available in primary and DR sites. In each case, both the VMware and Citrix DR solutions rely on SAN-based replication technologies to keep VM files and configuration data current at a backup location. VMware offers Site Recovery Manager, an add-on product that is essentially a workflow engine to orchestrate the DR of systems virtualized with vSphere. VMware licenses Site Recovery Manager on a per-protected-VM model. Unlike vSphere, XenServer includes Integrated Disaster Recovery in the Platinum Edition, with the license covering an unlimited number of VMs. Integrated Disaster Recovery simplifies configuration of VM recovery across primary and secondary sites, allowing failover to a secondary site and fail back once the primary site has recovered.

Memory Optimization and Dynamic Memory Control

Memory optimization involves guaranteeing minimum levels of memory as well as efficiently reclaiming unused memory and allocating it to VMs that require it. Both XenServer and vSphere offer memory optimization technologies as part of their virtualization platforms. vSphere leverages guest ballooning, page sharing and memory compression as the primary memory optimization technologies and will over-commit the physical memory on the vSphere host if required. In the event over-commitment occurs, the hypervisor will swap out host memory to disk with a corresponding performance impact. Memory compression seeks to minimize this performance impact by reducing the time to store and retrieve the memory from disk.

Page sharing is a memory optimization technique designed under the assumption that core operating system modules used in multiple VMs are identical and can be shared with a boost in memory available to applications. While this assumption is valid in older operating systems, many modern operating systems have implemented security and performance optimizations that minimize the viability of page sharing.

Because of these security and performance concerns, dynamic memory optimization in XenServer is based on the concept of a guest balloon. Dynamic memory control allows administrators to define memory boundaries in which a workload can meet its service level agreements without allocating excessive memory or incurring performance penalties due to swap. XenServer dynamic memory control allows the host memory to be over-subscribed, but never overcommitted. Dynamic memory control is available in all premium versions of XenServer and seamlessly works with critical components such as XenMotion and workload balancing.

Virtual Network Switch

A virtual network switch is a logical switching fabric that is built into the virtual infrastructure and enables management of virtualized network and security profiles, as well as virtual machine configurations, as they migrate across physical hosts. Distributed virtual network switching allows a multi-tenant, highly secure and extremely flexible network fabric to be created, enabling customers to move beyond server consolidation and into dynamic resource allocation. XenServer and vSphere both offer distributed virtual switches.

In XenServer the distributed virtual switch (DVS) is part of the free version and the virtual switch controller, which allows network management, grouping and traffic shaping, is included in all premium versions. XenServer DVS controller supports critical network operations tasks such as monitoring VM traffic (RSPAN), network access and performance monitoring (NetFlow) and access control restrictions (ACLs).

In vSphere, the virtual distributed switch (vDS) requires vSphere Enterprise Plus. Customers with lower editions can only leverage what VMware refers to as the “standard switch”. The standard switch lacks critical functionality like centralized network management. This gap requires virtualization administrators to manually configure network settings on each host in a cluster manually.

To leverage the full capabilities of VMware vDS, an add-on component from Cisco is required. The Cisco Nexus 1000v Switch is made up of two components, a Virtual Ethernet Module which runs in ESX/ESXi host, and the external Virtual Supervisor Module (VSM) to manage the VEMs. The Nexus 1000v provides a rich switching fabric to vSphere Enterprise Plus customers.

System Maintenance

Both XenServer and vSphere require some regular maintenance to apply software updates and patches. Because both offer live migration features, patching and updates to the hypervisor can be performed without incurring downtime for VMs. Both vSphere and XenServer offer the ability to perform pool wide upgrades of the hypervisor. In vSphere this functionality is contained within the Update Manager, and within XenServer this is implemented using the rolling pool upgrade wizard.

At a Glance: XenServer and VMware

	Citrix XenServer 6	VMware vSphere 5
Pricing model	Per server: No restriction on processors or cores per processor	Per CPU socket and configured virtual memory(vRAM): Penalizes use of more-powerful servers and higher consolidation ratios
Multi-core processor support	No restrictions	No restrictions
Bare-metal deployment	Free	Yes
P2V and V2V migration tools	Free	Yes
Multi-server management	Free	vCenter Server is $6,000-$8,000 extra.
Resource pools, shared storage	Free	Yes: Standard Edition and higher
VM snapshots	Free	Yes: Standard Edition and higher
Real-time performance monitoring	Free	Yes: Standard Edition and higher
Live motion, VM backup enablement	Free	Yes: Standard Edition and higher
VM high availability (HA)	Yes, Advanced Edition	Yes: Standard Edition and higher
HA of core management services	Yes, Advanced Edition	vCenter Server Heartbeat is $12,493 extra.
Historical performance monitoring	Yes, Advanced Edition	Yes: Standard Edition and higher
Administrator alerts	Yes, Advanced Edition	Yes: Standard Edition and higher
Dynamic memory configuration	Yes, Advanced Edition	Yes; Standard Edition and higher
Distributed switch	Yes, Advanced Edition	Yes; Enterprise Plus Edition only. Full capabilities require separate purchase of Cisco Nexus 1000v at $869 per processor.
StorageLink	Yes, Enterprise Edition	Nothing comparable
Provisioning services for VMs	Yes, Enterprise Edition	Nothing comparable
Workload balancing	Yes, Enterprise Edition	Yes: Enterprise Edition and higher
Role-based administration	Yes, Enterprise Edition	Yes; Standard Edition and higher
Disaster recovery enablement	Yes, Platinum Edition	Site Recovery Manager: $495 per protected VM
VM fault tolerance	Products from Marathon Technologies and Stratus	Yes: Enterprise Edition and higher (for VMs with 1 vCPU)

Licensing and Pricing

As this paper has shown, XenServer and vSphere both provide enterprise virtualization and management functionality. In addition XenServer, by leveraging the inherent benefits of paravirtualization and virtualization assist technologies, provides increased workload performance and density. The final area of comparison is licensing and pricing and how they translate into business value.

In this arena Citrix differs dramatically from VMware. Rather than a per-processor, configured virtual memory model, XenServer is licensed per host server and support is offered on a per-instance basis. This approach provides additional value for the customer looking to maximize a virtualization budget and stands in stark contrast to the per-processor, configured memory and VM licensing path, as the following examples show.

Sample Pricing Comparisons

The following example shows cost comparisons of XenServer and vSphere for the consolidation of physical servers.

Example 1: Consolidation of 120 physical servers with 2GB RAM to 15, using two-processor servers and an 8:1 consolidation ratio.

XenServer 6 Advanced Edition	**VMware vSphere 5 Standard Edition**
Capabilities: Multi-server management, resource pools, XenMotion, HA, dynamic memory	Capabilities: Multi-server management, resource pools, vMotion, HA, memory optimization
15x XenServer Advanced Edition: $15,000 (Includes 1 year of Citrix Subscription Advantage™) 24x7 technical support: $3,000 x 3 years = $9,000 SA renewal: $130 / host x 2 years = $3900	vSphere Standard, 3-year production support = $54,560 (licensing requires 30 processor licenses) vCenter Server, 3-year production support = $8,180 vCenter Server Heartbeat, 3-year production support = $16,366
Total three-year cost: $27,900	Total three-year cost: $79,106
XenServer Savings: 65%	

Example 2: Consolidation of 120 physical servers with 4GB RAM to 8, using two-processor servers and a 15:1 consolidation ratio.

XenServer 6 Advanced Edition	**VMware vSphere 5 Standard Edition**
Capabilities: Multi-server management, resource pools, XenMotion, HA, dynamic memory	Capabilities: Multi-server management, resource pools, vMotion, HA, memory optimization
8x XenServer Advanced Edition: $8,000 (Includes 1 year of Citrix Subscription Advantage™) 24x7 technical support: $3,000 x 3 years = $9,000 SA renewal: $130 / host x 2 years = $2,080	vSphere Standard, 3-year production support = $36,373 (licensing requires 20 vRAM licenses) vCenter Server, 3-year production support = $8,180 vCenter Server Heartbeat, 3-year production support = $16,366
Total three-year cost: $19,080	Total three-year cost: $60,919
XenServer Savings: 69%	

Example 3: Consolidation of 120 physical servers with 8GB RAM to 12, using four-processor servers and a 10:1 consolidation ratio.

XenServer 6 Enterprise Edition	VMware vSphere 5 Enterprise Plus
Capabilities: Multi-server management, resource pools, XenMotion, dynamic memory, HA, WLB, distributed network switch	Capabilities: Multi-server management, resource pools, vMotion, memory optimization, HA, DRS, distributed switch, Cisco Nexus 1000v
12x XenServer Enterprise Edition: $30,000 (Includes 1 year of Subscription Advantage) Technical Support: $3,000 x 3 years = $9,000 SA renewal: $325 / host x 2 years = $7,800	vSphere Enterprise Plus, 3-year production support = $320,035 (licensing requires 48 processor licenses) vCenter Server, 3-year production support = $8,180 vCenter Server Heartbeat, 3-year production support = $16,366
Total three-year cost: $46,800	Total three-year cost: $344,581
XenServer Savings: 86%	

Example 4: Multiple site consolidation of 240 physical servers with 8GB RAM to 12, using four-processor servers and a 20:1 consolidation ratio; includes disaster recovery protection of 100 critical servers requiring 5 backup virtualization hosts.

XenServer 6 Platinum Edition	VMware vSphere 5 Enterprise Plus
Capabilities: Multi-server management, resource pools, XenMotion, dynamic memory, HA, WLB, distributed network switch, Site Recovery Manager	Capabilities: Multi-server management, resource pools, vMotion, memory optimization, HA, DRS, distributed switch, Cisco Nexus 1000v, Site Recovery Manager
17x XenServer Platinum Edition: $85,000 (Includes 1 year of Subscription Advantage) Technical Support: $3,000 x 3 years = $9,000 SA renewal: $675 / host x 2 years = $22,950	vSphere Enterprise Plus, 3-year production support = $320,035 (licensing requires 48 vRAM licenses) 2x vCenter Server, 3-year production support = $16,360 2x vCenter Server Heartbeat, 3-year production support = $32,732 100x vCenter Site Recovery Manager = $73,693
Total three-year cost: $116,950	Total three-year cost: $442,820
XenServer Savings: 74%	

*Pricing from VMware online store as of November, 2011

Conclusion

Just a few years ago, customers looking for server virtualization solutions had limited options. The server virtualization market has entered a new phase in which underlying innovations from Intel and others have enabled the development of powerful and capable solutions such as Citrix XenServer. For customers, XenServer offers a clear alternative to vSphere. With XenServer, Citrix not only provides a feature-rich virtualization platform that meets or exceeds all the functionality of vSphere but, as this white paper has shown, XenServer surpasses vSphere in a number of areas including its open architecture, performance, storage integration and business value.

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is a leading provider of virtual computing solutions that help companies deliver IT as an on-demand service. Founded in 1989, Citrix combines virtualization, networking, and cloud computing technologies into a full portfolio of products that enable virtual work styles for users and virtual datacenters for IT. More than 230,000 organizations worldwide rely on Citrix to help them build simpler and more cost-effective IT environments. Citrix partners with over 10,000 companies in more than 100 countries. Annual revenue in 2010 was $1.87 billion.