SUN AND VMWARE SOLUTIONS GUIDE FOR WINDOWS SERVER CONSOLIDATION

White Paper
June 2009

Abstract

Designing a Windows server consolidation initiative creates a new set of challenges for the storage infrastructure — storage solutions must offer simplicity, performance, availability, and ease of management. To this end, Sun offers a range of simple, high-performance, easy-to-deploy, low-cost, and innovative storage solutions for Windows consolidation efforts. When combined with VMware vSphere 4.0 (or earlier versions of VMware software), Sun provides a solution for Windows virtualization that leverages industry-standard file access protocols and the extensive features of Sun™ Storage 7000 Unified Storage Systems.

Table of Contents

Simplifying storage in VMware deployments

Today, more than ever, IT professionals are deploying virtualization technologies to consolidate and simplify the infrastructure for delivering a broad spectrum of Microsoft Windows services. Typically, virtualization initiatives seek to reduce datacenter costs by consolidating multiple physical servers, lowering energy expenses, and reducing administrative overhead. In addition, when designing virtualized solutions, system architects strive to improve the resiliency of Windows-based IT infrastructures and provide high levels of application and data availability.

VMware offers a proven virtualization infrastructure that is implemented by over 130,000 customers to date. The latest release of its flagship software product — recently renamed VMware vSphere 4.0 — is designed to provide a full virtualization solution that supports cloud computing. Cloud computing refers to an architectural abstraction in which applications and IT infrastructure resources are pooled to deliver application services as needed via the network (which is commonly depicted in architectural diagrams as a cloud). Implementing a virtualization infrastructure using VMware aims to optimize resource utilization, reduce IT acquisition and operating costs, and enable greater control and flexibility in managing the overall IT infrastructure.

Storage solutions that optimize simplicity, performance, availability, and ease of management provide synergy with VMware when deploying a virtualization infrastructure. Sun offers innovative storage products and extensive data services in the Sun™ Storage 7000 Unified Storage Systems, enabling easy-to-manage, low-cost storage resources that span a range of capabilities. This family of network-attached storage appliances takes advantage of industry-standard file system protocols, including NFS (Network File System) for virtual machine storage and CIFS (Common Internet File System) for Windows virtual machine access. Because of the simplicity offered by the implementation of these industry-standard protocols, Sun Storage 7000 Unified Storage Systems eliminate much of the complexity typically associated with common storage configuration and management tasks. In addition, these appliances deliver high I/O performance with low latency in configurations that feature both read and write caching functionality. Sun Storage 7000 Unified Storage Systems also offer a full range of data services — compression, backups, replication, and real-time performance analytics — that complement the appliances' high performance and ease of operation.

Sample reference architectures for VMware deployments

To meet a wide spectrum of workloads, this paper proposes four reference architectures targeted at VMware deployments. These sample reference architectures provide a starting point for implementing virtualization solutions, and can be easily customized and extended to meet specific business needs. System

architects can scale storage, server, networking, and other architectural components as needed to meet site requirements or to accommodate anticipated growth in capacity over time.

The reference architectures presented in this paper are based on Sun Storage 7000 Unified Storage Systems and Sun AMD® Opteron or Intel® Xeon® processor-based servers. They are designed to speed time-to-production and to achieve a new dimension of efficiency characterized by simplicity, high-performance, powerful analytics, and low cost. Compared to the complexity that often results from using SAN storage offerings, Sun Storage 7000 Unified Storage Systems enable straightforward, file-based virtual machine storage through NFS. When it comes to NFS technology, no vendor has more experience than Sun, who invented NFS initially and revolutionized file sharing technology in the industry. Over the last 25 years, NFS has become a proven standard that is ideal for facilitating storage within a virtualization infrastructure.

In addition to NFS protocols, Sun Storage 7000 Unified Storage Systems include native CIFS support that simplifies data access by Windows virtual machines. Sun Storage 7000 Unified Storage Systems incorporate an in-kernel CIFS/SMB implementation in the OpenSolaris™ Operating System. A rich set of permission and identity mapping capabilities allows transparent file access — even to files on the same share — from Windows or UNIX virtual machines. By simplifying Windows and UNIX interoperability and seamlessly mapping file system credentials, the storage appliance can integrate more easily into existing heterogeneous environments. In this way the appliance speeds deployment and eases operations, lowering administrative overhead and reducing cost.

"Storage virtualization" versus "storage in server virtualization"

It is important to clarify terminology with regards to storage resources and virtualization. This paper specifically addresses how to address storage requirements typical in virtualization initiatives with VMware infrastructure software. In this context, "storage in server virtualization" refers to storage used to support applications and operating systems that are virtualized and encapsulated in a virtual machine (the virtualized application stack) and related data (VMDK disks). This storage is a fundamental element of a consolidation scenario where x64 servers run VMware software. Careful storage planning to support the virtualized application stack and related data is necessary prior to deploying a consolidated Windows server environment.

In contrast, the term "storage virtualization" refers to the pooling of multiple physical storage resources into what appears as a single storage resource that is centrally managed — a topic that falls outside the scope of this paper. Although Sun offers storage virtualization solutions, this paper does not discuss them.

Instead, this paper focuses on storage solutions that can be deployed to support server virtualization initiatives. In particular, it centers on the simplicity of using the Sun Storage 7000 Unified Storage Systems and Sun x64 servers in conjunction with the VMware software. Although VMware has built-in support for standard block protocols — Fibre Channel, iSCSI, SCSI, and SAS — as well as NFS protocols for virtual machine storage, these reference architectures recommend using NFS on Sun Storage 7000 Unified Storage Systems because of the overriding simplicity that this solution offers. Along with the simplicity of file-based access through NFS, these appliances provide a broad spectrum of capabilities with respect to data services, performance, capacity, high availability, and storage analytics.

Virtualization solutions from Sun Microsystems and VMware

Sun supplies servers and storage systems that are ideally suited for virtualization and certified to work well with VMware vSphere 4.0 software (and earlier VMware software versions). An on-line VMware Compatibility Guide lists all Sun products currently certified for VMware — simply search for the key words "Sun Microsystems" at www.vmware.com/resources/compatibility, or see www.sun.com/software/vmware/certification.jsp. Although many Sun products are certified for VMware support, the following Sun server and storage models are particularly well-suited for VMware infrastructure solutions:

- Sun x64 rack-mounted servers with AMD or Intel processors, including:
  - Sun Fire X4170 Server
  - Sun Fire X4270 Server
  - Sun Fire X4450 Server
  - Sun Fire X4440 Server
  - Sun Fire X4600 M2 Server
- Sun x64 blade servers with AMD or Intel processor blades (such as Sun Blade 6000 and Sun X6240 or Sun X6270 Processor Blades)
- Sun Storage 7000 Series Unified Storage Systems, including:
  - Sun Storage 7110 Unified Storage System
  - Sun Storage 7210 Unified Storage System
  - Sun Storage 7310 Unified Storage System
  - Sun Storage 7410 Unified Storage System

Since Sun resells and supports VMware software, Sun provides a single focal point for customers who seek integrated solutions encompassing infrastructure storage, server hardware, VMware software, and support. This end-to-end solution capability is provided by very few vendors and is noticeably absent from the portfolios of storage-only vendors.

Sun and VMware continue to pursue compatible technology directions with respect to virtualization and related technologies, focusing on delivering solutions that help to scale workloads, conserve power, improve business continuity, and enhance performance observability — directions that ultimately help customers to simplify their IT operations and reduce cost.

Given recent enhancements to VMware vSphere 4.0, Sun x64 servers and Sun Storage 7000 Unified Storage Systems are especially well-suited as underlying compute and storage platforms for virtualization. With VMware vSphere 4.0, a virtual machine can now scale to support up to 64 cores and 512GB of physical memory. This allows the number of virtualized servers to scale more easily in support of more demanding workloads, closely complementing the good scalability, large memory capacities, and number of processor cores available in Sun's portfolio of x64 servers. To ramp storage performance and reduce I/O latency, certain Sun Storage 7000 Unified Storage System models can be configured with enterprise-quality solid state devices (SSDs) to cache read and write operations. For example, the Sun Storage 7410 system can support up to 600 GB of read cache and 288 GB of write cache, which can help to cache virtual machine execution and accelerate read and write performance for I/O-intensive applications. The reference architectures offer a spectrum of storage capacities and capabilities, reflecting the scalability of both the Sun x64 servers and Sun Storage 7000 Unified Storage Systems.

Reference architectures

For the proposed reference architectures, Table 1 summarizes baseline configurations, which are further discussed in the remainder of this paper.

Table 1. Baseline configurations

Entry-level sandbox	**Small: Test & development**	**Medium: Pre-production, staging**	**Large: Production**
Ideal for demonstrations or training staff new to VMware	Ideal for developers and unit testing	Ideal for architects as a staging environment (without full-blown production data)	Ideal for high data availability, capacity and performance requirements
Sun Fire X4270 server, fully populated	(2) Sun Fire X4170 servers	(2) Sun Fire X4170 servers	(2) Sun Fire X4170 servers
RAID 5 SAS	Sun Storage 7110 Unified Storage System with mirroring, striping, double parity, Hybrid Storage SAS disks	Sun Storage 7210 Unified Storage System with mirroring, striping, double parity, Hybrid Storage SATA-II disks and SSDs w/ flash acceleration	Sun Storage 7410 Unified Storage System with mirroring, striping, double parity, Hybrid Storage SATA II-disks and SSDs w/ flash acceleration; options for high availability
2TB internal raw storage	2 TB or 4 TB internal raw storage	11.5 - 142 TB raw storage	44 TB - 288 TB raw storage (with a future capacity of 576 TB)

Entry-level sandbox

The entry-level configuration with SAS storage allows customers to investigate how applications perform in a virtualization environment. The goals behind this configuration are delivering low cost, simplicity, and ease of deployment. For environments that require more advanced virtualization features, a medium, large, or customized configuration is more appropriate.

For the entry-level configuration, the VMware server is configured with a minimum of 24 GB of memory and an Intel® Xeon® or AMD® Opteron processor. Local storage is often suitable for small capacity requirements, but direct-attached storage can be implemented for larger capacity requirements.

Figure 1 shows a typical entry-level configuration using a Sun Fire X4270 server connected to a network switch and then to the rest of a customer network. Based on state-of-the-art Intel® Xeon® Processor 5500 Series technology (code-named "Nehalem"), the Sun Fire X4270 server comes standard with four Gigabit Ethernet ports, allowing one port to be dedicated to management and backup and the remaining ports to service the virtual machines. VMware software is installed on the Sun Fire X4270 server and can host a number of virtual machines. The virtualization software can be installed on the first two disks, which can be mirrored for greater resiliency. The virtual machines can then use the remaining 14 disks in a RAID 5 pool. When RAID 5 data is striped across the 14 disk drives with distributed parity, it enables approximately 1.7 TB of usable storage.

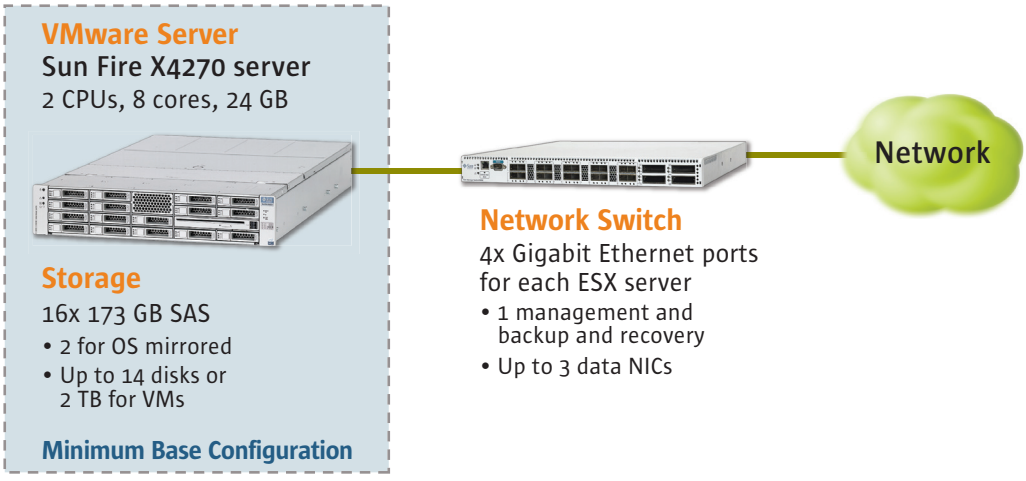

Figure 1. Typical entry-level sandbox configuration.

Small configuration – test and development

The small configuration is designed to support at least 10 virtualized servers in environments without significant capacity storage demands. This reference architecture focuses on delivering enhanced uptime, maintenance scheduling, resource scheduling, and manageability. Although any Intel or AMD server can be deployed in this configuration, a Sun Fire X4170 server is selected in this reference architecture to support the VMware server (Figure 2). In this role, the Sun Fire X4170 server is configured with 72 GB of memory and two Intel Xeon processors.

Additional servers can be deployed in this configuration based on the need to scale the number of virtual machines and to offer server redundancy. Generally, it is recommended that multiple VMware servers are deployed so that not all virtual machines are configured on one physical machine. With multiple servers, virtual machines can be migrated to other servers using VMotion, which accommodates planned and unplanned downtimes. All VMware servers can be managed with VMware vCenter Server (formerly VMware VirtualCenter). Each server is attached to shared NFS storage, which enables additional VMware features such as VMotion, High Availability, Distributed Resource Scheduling, and Storage VMotion.

Figure 2 illustrates a typical small configuration but does not reflect the extensive range of all possible configurations. Additional VMware servers and storage expansion trays can be added to this configuration. VMware vCenter server provides features to manage all VMware servers, the shared storage, and the license server. Either an existing Microsoft SQLServer or a new server can be installed for the database. Each of the VMware servers has one dual-port and one quad-port networking card installed to provide a total of 10 network ports (the server platform itself features four on-board ports). A dedicated network port is assigned for management, VMotion, and backup and recovery, along with two ports allocated for storage. The five remaining network ports are used for data access for the virtual machines.

To support storage requirements in this configuration, the Sun Storage 7110 Unified Storage System can be purchased with either 2.0 TB or 4.0 TB of capacity in an energy-efficient 2U package. Storage disk drives can be configured with mirroring, striping, or double parity. The configuration can be set up to take unlimited snapshots (depending on capacity) and provide as many volumes as needed (again, subject to available capacity).

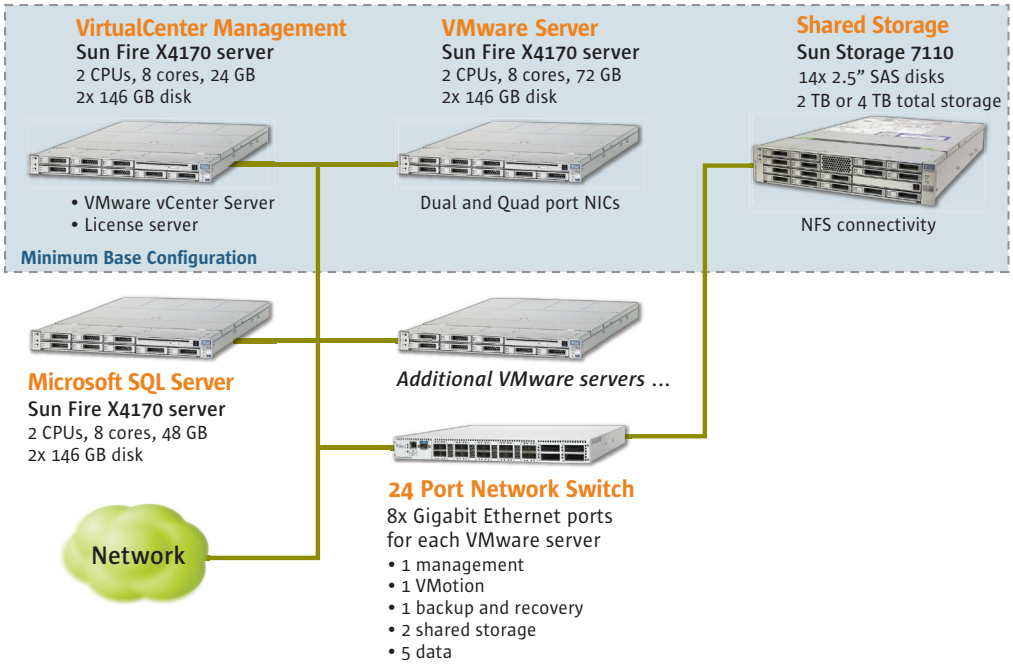

Figure 2. Typical small configuration.

Medium configuration – pre-production, staging

The medium configuration is designed to meet the needs of business customers with approximately 50 or more servers and up to 144 TB. The medium configuration focuses on delivering enhanced uptime, maintenance scheduling, resource scheduling, and manageability.

Although any Intel or AMD server can be deployed in this configuration, the Sun Fire X4170 server runs VMware software and is configured with 72GB of memory and two high-performance Intel Xeon processors. A number of additional servers can be used as needed to meet requirements for the number of virtual machines and to satisfy redundancy requirements. Additional servers are recommended so that not all virtual machines are configured on one physical machine. With multiple servers, virtual machines can be migrated to other servers using VMotion to address planned and unplanned downtimes. The VMware servers are all managed with VMware vCenter server. Each server is attached to shared storage that is accessed through NFS, enabling additional VMware features (such as VMotion, High Availability, Distributed Resource Scheduling, and Storage VMotion) and supporting a variety of Microsoft applications (such as SQLServer, Exchange, Sharepoint, Great Plains, home directories, compliance data, Office File proliferation, etc.).

Figure 3 illustrates a typical medium configuration but does not reflect the full range of possible configurations. Additional VMware servers and storage expansion trays can be added to expand compute and storage resources in this configuration.

This configuration supplies network-attached storage through NFS for virtual machines. VMware VCenter server provides features needed to manage all VMware servers, the shared storage, and the license server. Either an existing Microsoft SQLServer or a new server can be installed for the database. Each of the VMware servers has one dual-port and one quad-port networking card installed, providing a total of 10 network ports (the Sun server has four on-board ports). A dedicated network port can be issued for management, VMotion, and backup and recovery, along with two ports assigned to storage. The five remaining network ports can be used for data access by the virtual machines.

The Sun Storage 7210 Unified Storage System supports 11.5 TB to 48 TB in an energy-efficient 4U package and can be augmented with two additional 4U expansion arrays to achieve a maximum capacity of 142 TB. If greater storage capacity is needed for staging the pre-production environment, the Sun Storage 7410 Unified Storage System can be substituted, enabling a maximum capacity of 288 TB (with a future capacity of 576 TB). All Sun Storage 7000 Unified Storage Systems feature an easy-to-use GUI-based installation front-end and are easily configured in a matter of minutes. Storage disk drives can be configured with striping, mirroring, and double parity. The storage appliance can also be configured to take unlimited snapshots (depending on capacity) and to provide as many volumes as needed (also depending on available capacity).

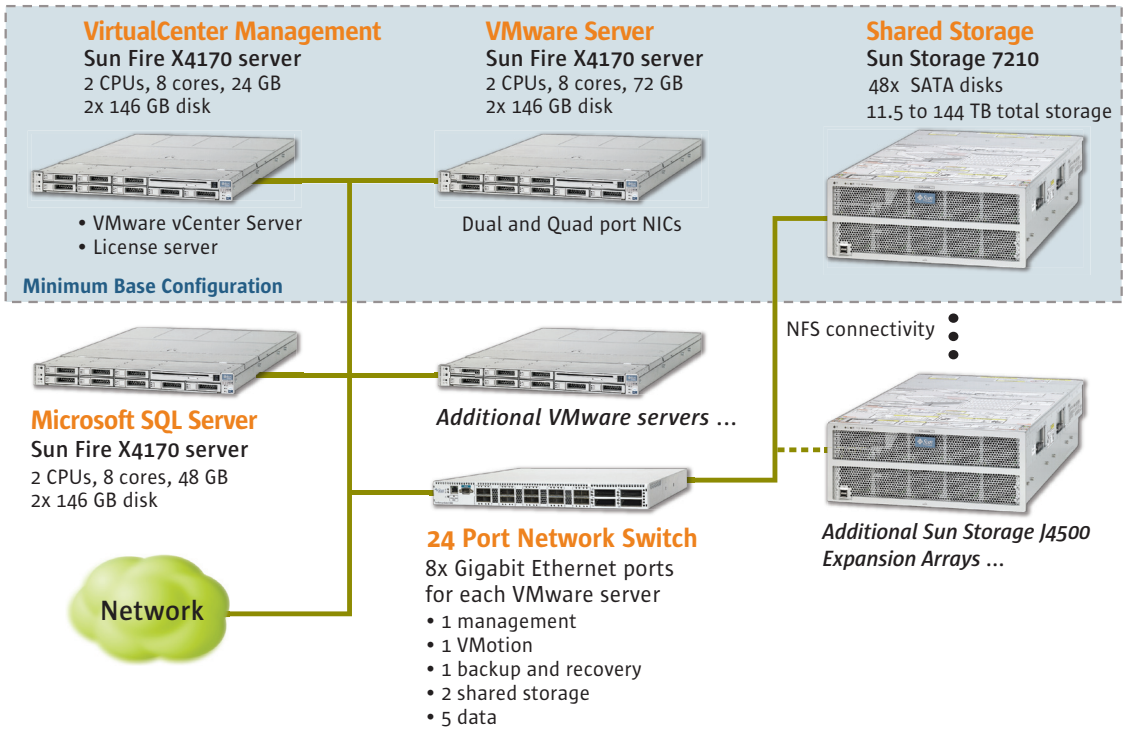

Figure 3. Typical medium configuration.

Large configuration – production

The large configuration can meet the needs of business customers that require consolidation of approximately 100 or more servers. This configuration also focuses on high performance, uptime, maintenance scheduling, resource scheduling, and manageability.

The high performance afforded by this reference architecture configuration is a direct result of solid state drives (SSDs) in the Sun Storage 7410 Unified Storage System. reaed- and write-optimized SSDs help to accelerate VM cache performance because of flash technology's low latency. Hybrid Storage Pool functionality automatically optimizes data placement to easily take advantage of this capability.

A Sun Storage 7410 Unified Storage System starts with 24 TB of capacity in an energy-efficient 4U package and can be further expanded up to 288 TB (or up to 576 TB in the future) using up to 11 additional 4U storage expansion arrays. Storage drives can be configured with striping, mirroring, and double parity. The storage system can be set up to take unlimited snapshots (depending on capacity) and to provide as many volumes as needed (also depending on capacity).

Either a blade-based or rack-mounted Sun x64 server can be used to run VMware services in the large configuration. As an example, either a series of AMD or Intel-processor-based servers (such as multiple Intel-based Sun Fire X4170 servers) or a Sun Blade 6000 server (with a series of AMD or Intel processor blades) can function as the VMware server. The broad portfolio of Sun x64 servers incorporates processor technologies recently introduced by both Intel and AMD, and are generally available in both rack-mounted and blade-based form factors.

As with the other reference architecture configurations, additional VMware servers can be added depending on the number of virtual machines needed and to enhance redundancy. Additional physical servers are recommended so that not all virtual machines are configured on a single physical machine. With multiple servers, virtual machines can be migrated with VMotion to other servers to accommodate both planned and unplanned downtime. VMware servers are all managed with VMware vCenter server. Each server is attached to shared storage that is accessed via NFS, enabling additional VMware features (such as VMotion, High Availability, Distributed Resource Scheduling, and Storage VMotion) and supporting a spectrum of applications (such as SQLServer, Exchange, Sharepoint, Great Plains, Office File proliferation, compliance data, data marts, DB cubes, etc.).

Figure 4 illustrates a typical large configuration but does not reflect all possible configurations. Additional virtualization servers and storage expansion trays can be added. Network-attached, file-based VM storage in this configuration (using NFS) greatly simplifies setup and management. VMware vCenter server provides capabilities to manage all virtualization servers and shared storage, as well as the

license server. Either an existing Microsoft SQLServer or a new server can be installed for the database. Each of the VMware servers has one dual port and one quad port network cards installed for a total of 10 network ports (the system natively has four ports). A dedicated network port is issued for management, VMotion, and backup and recovery, as well as two ports for VM access. As shown in Figure 4, 10 Gigabit Ethernet networking can be used to support high-bandwidth data networking for the virtual machines.

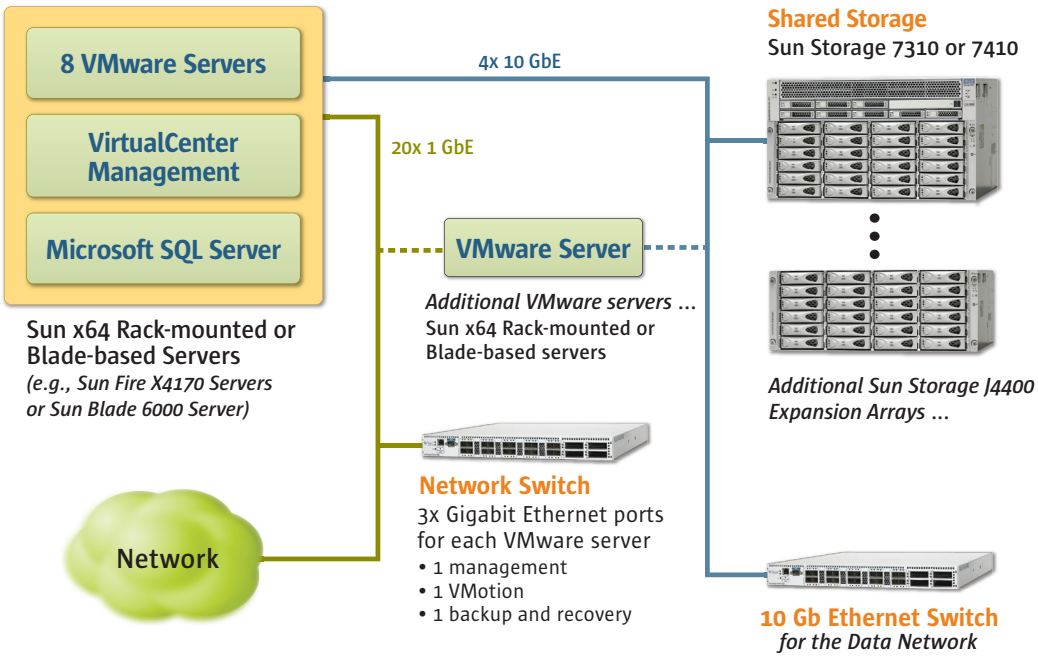

Figure 4. Typical large configuration.

Configuring high-availability storage

Clustering with Sun Storage 7310C or 7410C Unified Storage Systems provides an optimal solution for sites with stringent availability requirements. With active/passive clustered configurations of Sun Storage 7310C or 7410C Unified Storage Systems, a second server node in the appliance functions as a hot spare — providing failover and protection in the event of the loss of the primary server node. Active/active configurations are also supported.

Figure 5 depicts a typical enterprise configuration with a clustered Sun Storage 7410C Unified Storage System. As in the large configuration, this configuration is designed to meet the needs of business customers with approximately 100 or more servers and focuses on high performance, uptime, maintenance scheduling, resource scheduling, and manageability. The high-performance storage capabilities in this configuration result from using solid state drives and flash acceleration in the Sun Storage 7410C Unified Storage System. These technologies enable rapid write

capabilities for fast data placement on the storage pool and help to optimize I/O rates. This boosts VM cache performance with NFS, especially because of the low latency and high performance possible with flash technology.

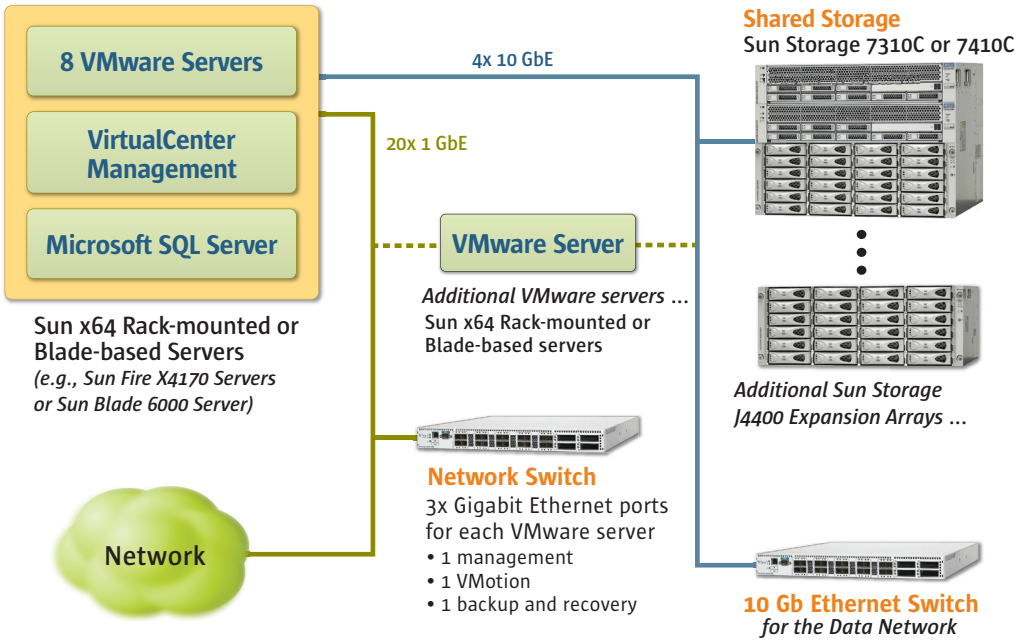

Figure 5. Typical enterprise configuration with high availability (using a clustered Sun Storage 7410C Unified Storage System).

A Sun Storage 7410C Unified Storage System is available with up to 288 TB (and up to 576 TB in the future), depending on customer needs. Up to 12 Sun Storage J4400 expansion arrays are used to maximize capacity. The Sun Storage J4400 array offers enterprise-class JBOD (Just a bunch of disk) capabilities with dual controllers for robust connectivity. Each expansion array houses 24 3.5-inch SATA-II devices within a compact 4U form-factor. Like the Sun Storage 7410C Unified Storage System, the Sun Storage 7310 Unified Storage System is also available in a clustered configuration for high availability. At a lower entry-level price point, the Sun Storage 7310 System is designed to meet lower capacity requirements — it supports up to 96 TB in up to four additional expansion arrays.

For both models, storage disk drives can be configured with mirroring, striping, and double parity. The storage system can also be set up to take unlimited snapshots (depending on capacity) and provide as many volumes as needed (also depending on available capacity).

Either blade-based or rack-mounted Sun x64 servers can be used to run VMware services in a large, high availability configuration. As an example, either a series of AMD or Intel-processor-based servers (such as multiple Intel-based Sun Fire X4170 servers) or a Sun Blade 6000 server (with ten Sun X6270 processor blades) can be configured to run the VMware software.

Multiple servers can be used based on the number of total number of required virtual machines and the need for redundancy. Additional servers are recommended so that not all of the virtual machines are configured on a single physical server. With multiple servers, virtual machines can be migrated to other servers using VMotion to address planned or unplanned downtime. Each server is attached to shared storage that is accessed via NFS, enabling additional features such as VMotion, High Availability, Distributed Resource Scheduling, and Storage VMotion.

Figure 5 illustrates a typical large configuration but does not reflect all possible large high availability configurations. Additional virtualization servers and storage expansion trays can be added as needed. VMware vCenter server provides all features to manage all VMware servers, the shared storage, and the license server. Microsoft SQLServer is also installed on one of the servers.

If the Sun Blade 6000 server is used, all ten server blades have two 1 Gigabit Ethernet NICs. Two shared 10 Gigabit Ethernet ports per blade server can be installed using two Sun Blade 6000 Virtualized Multi-Fabric Network Express Modules. The two 1 Gigabit Ethernet network ports for each blade can be issued for management, VMotion, and backup and recovery. The remaining two 10 Gigabit Ethernet interfaces can be used for NFS services to the Sun Storage 7310C or 7410C Unified Storage System for the virtual machines, supplying plenty of bandwidth and cable aggregation for virtual machine data.

Advantages of file-based storage using NFS

In all of the reference architecture configurations, Sun Storage 7000 Unified Storage Systems offer significant advantages for virtual machine storage:

- *Simplicity.* In contrast to other, more complex storage solutions such as Fibre Channel-based SAN storage, network-attached storage using NFS takes only minutes to configure, accelerating time-to-deployment and speeding data access for virtual machines. The simplicity of NAS also translates into low cost since, unlike Fibre Channel-based solutions, no specialized switches or interfaces are required. Existing infrastructure TCP/IP networks and switches can be used to support storage access.
- *High-performance.* Built into Sun Storage 7000 Unified Storage Systems, Hybrid Storage Pool technology enables the pooling of DRAM, SSDs, and low-cost hard disk drives, optimizing performance according to workloads.

- *Observability.* DTrace Analytics provides in-depth and real-time monitoring of VM clients, VM disks, and I/O operations, which helps administrators to identify and resolve performance bottlenecks.
- *Interoperability.* To support file access for Windows virtual machines, Sun Storage 7000 Unified Storage Systems also support CIFS, allowing the appliance to seamlessly provide access for both Windows and UNIX clients, even within the same share.

Administrative simplicity

Sun Storage 7000 Unified Storage Systems take only minutes to configure for NFS or CIFS access, and Hybrid Storage Technology (described next) simplifies expansion when new devices are added. As shown in Figure 6, a browser-based GUI simplifies configuration tasks and helps to expedite time-to-deployment. When configuring CIFS storage for Windows virtual machines to use, the administrator identifies an existing Active Directory domain, defines a storage share, and establishes a set of identity mappings for Windows and UNIX client access. The GUI includes an option to automate virus scanning, enabling external virus detection software to scan and isolate quarantined files. Using these industry-standard protocols helps to speed initial deployment and streamline administrative tasks.

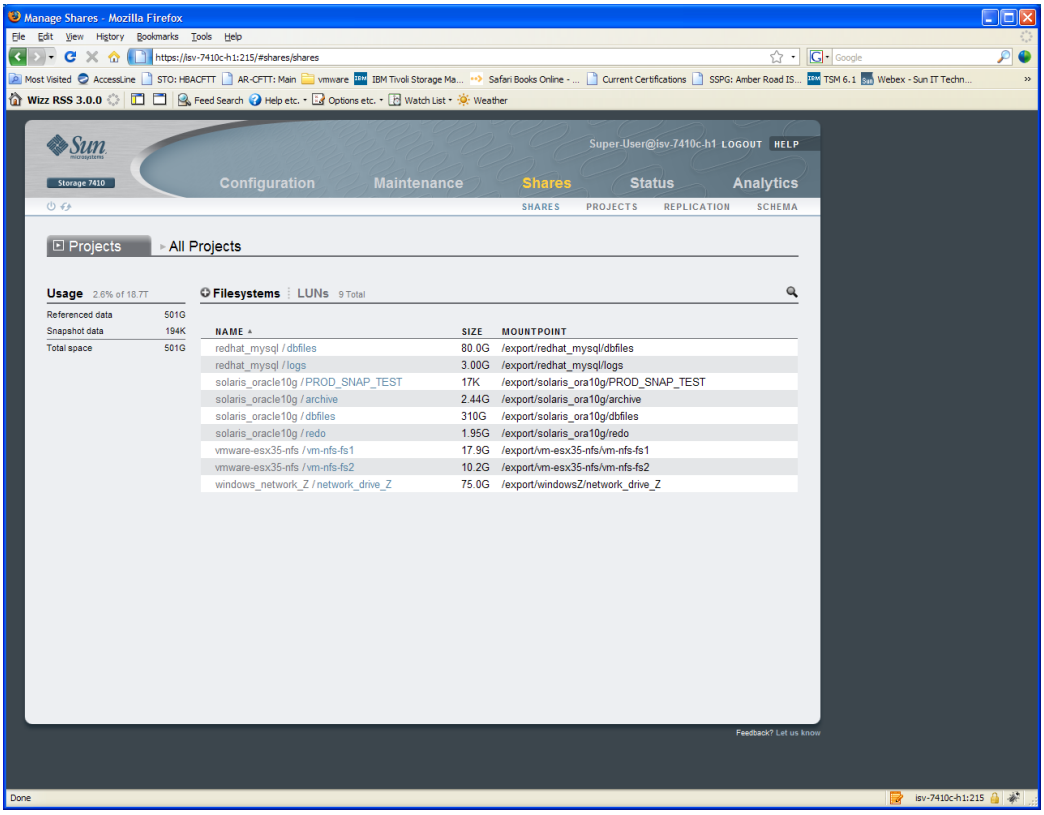

Figure 6. A Web-based GUI simplifies configuration and administrative tasks.

Thin provisioning

When VMware vSphere 4.0 creates a virtual machine on an NFS share, by default it efficiently allocates the virtual machine's disk file (its vdisk) using "thin provisioning." Thin provisioning assigns only the amount of space actually needed to hold the data rather than allocating the entire requested size for the vdisk. For example, suppose an administrator creates a Windows 2003 virtual machine with a vdisk size of 20 GB. A typical Windows 2003 installation takes approximately 2-3 GB of actual disk space. Without thin provisioning, the entire 20 GB of space would be allocated up front, resulting in 17-18 GB of space that does not immediately need to be allocated. In contrast, thin provisioning only allocates the 2-3 GB of disk space that is truly needed. Dynamic reservations and quotas for projects and shares give administrators additional flexibility to scale and manage storage capacities at optimal efficiency.

Flexible ZFS-based NAS offering

With its 128-bit foundation, ZFS effectively offers unlimited capacity and thin provisioning offers flexibility through the use of dynamic reservations and quotas for projects and shares. With these tools, administrators can both scale and manage storage allocated to various projects that share the storage appliance. As with all ZFS deployments, end-to-end checksumming of data and metadata is performed, eliminating the unpredictable effects of silent data corruption. Multiple data protection schemes are available, allowing administrators to make intentional trade-offs as they tune specific shares for capacity, availability, and performance (Figure 7).

Data Profile | Log Profile | Cache Profile

TYPE ▲	NSPF	AVAILABILITY	PERFORMANCE	CAPACITY	SIZE
Double parity RAID	No				5.37T
Double parity RAID, wide stripes	No				5.37T
Mirrored	No				3.36T
Single parity RAID, narrow stripes	No				4.03T
Striped	No				7.39T

Figure 7. ZFS configuration options allow administrators to tune individual shares for capacity, availability, and performance.

Hybrid Storage Pool technology

With Hybrid Storage Pool technology built into the Sun Storage 7000 Unified Storage Systems, virtual machines (and encapsulated applications) can transparently take advantage of the state-of-the-art enterprise solid state disk (SSD) technologies. Hybrid Storage Pools intelligently combine low-latency enterprise SSDs with high-capacity, highly reliable SATA HDDs. SSDs enable increased caching for ESX servers and virtual machines, which helps to accelerate I/O throughput. Hybrid Storage Pool

technology is designed to exceed the performance of Fibre Channel technologies without the additional management and administration complexity of a SAN, which can also suffer from extremely low utilization when used with disk shorting.

The ZFS architecture relies on these components (Figure 8) to implement Hybrid Storage Pools:

- The ZFS ARC (Adaptive Replacement Cache) is the main ZFS memory cache, stored in DRAM.
- The L2ARC (Level Two Adaptive Replacement Cache) extends the ARC into read-optimized SSDs (available in the Sun Storage 7310 and 7410 Unified Storage Systems) to provide a large read cache to accelerate reads.
- The ZIL (ZFS Intent Log) is transactional and uses write-based SSDs (available in the Sun Storage 7210, 7310, and 7410 Unified Storage Systems) to provide a large cache to accelerate writes.
- The disk storage pool consists of conventional disk drives. Note that high-performance, more costly disk drives are no longer strictly required to achieve high performance levels given the interposition of SSDs in a Hybrid Storage Pool.

Sophisticated filesystem algorithms in ZFS use the ARC in memory and the L2ARC on SSD to determine pre-fetch or data placement during sustained read operations. SSDs accelerate write throughput for ZFS synchronous write I/O operations, dramatically improving VM caching in VMware servers, which helps to boost write performance.

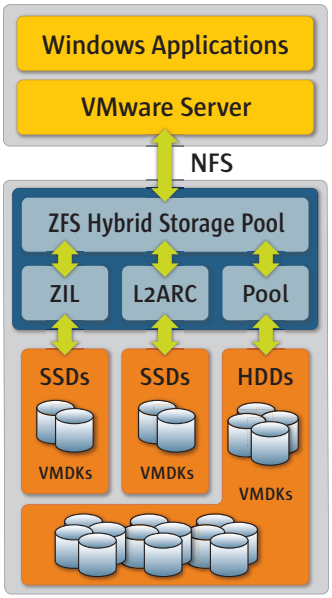

Figure 8. Hybrid storage pools optimize data placement to enhance I/O performance.

Real-time visibility with DTrace Analytics

After deployment into production, applications can sometimes exhibit throughput issues that were not exposed during initial testing. DTrace Analytics — a storage workload analysis and diagnostic technology built into the Sun Storage 7000 Unified Storage Systems — can help to isolate problems, improve VM storage performance, and enhance data availability.

As shown in Figure 9, DTrace Analytics provides both a top-down and bottom-up detailed view of storage workloads. It includes the ability to drill down into the detail of I/O throughput and VM storage workloads for purposes of isolating and removing performance bottlenecks, which can help to improve Microsoft application performance. Many key performance metrics — such as CPU and memory utilization, cache hits and misses as well as disk, file, and filesystem I/O operations per second — are captured and displayed in real time for each VMware virtual machine and its executing Microsoft applications. As Figure 9 shows, when a VM with a data warehousing application starts, the number of NFS operations per second performed by other VMs suddenly decreases.

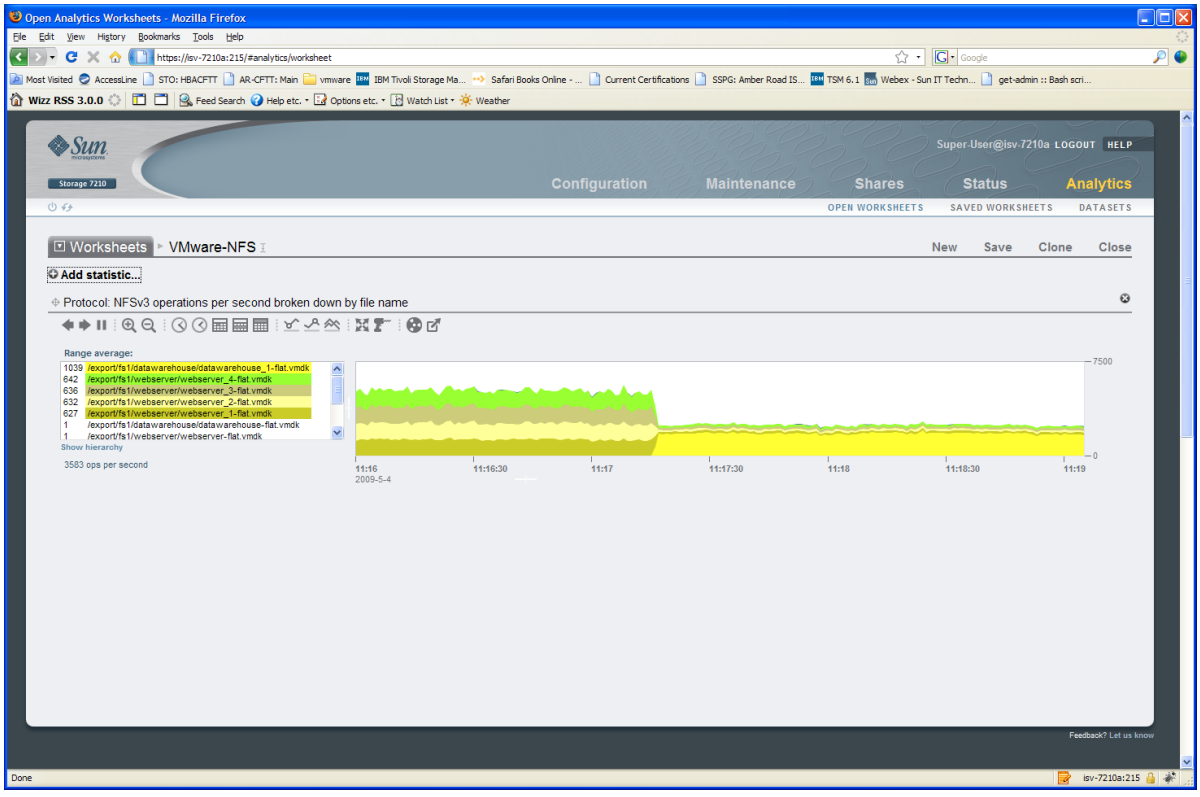

Figure 9. DTrace Analytics capabilities enable easy observability into resource use, improving storage monitoring and troubleshooting.

Data services

In conjunction with Sun Storage 7000 Unified Storage Systems, the storage infrastructure for virtualization can benefit from the broad set of available data services. Since both NFS and CIFS rely on underlying ZFS technology used in these storage appliances, these filesystems can take advantage of a rich set of available data services, including:

- *Network Data Management Protocol (NDMP) appliance backups.* Sun Storage 7000 Unified Storage Systems support both NDMP v3 and NDMP v4, and NDMP activities and statistics are available through the administrative interface. Sun StorageTek™ tape drives and virtual tape libraries (VTLs) are also supported.
- *Appliance-to-appliance remote replication.* Sun Storage 7000 Unified Storage Systems enable secure, asynchronous replication of storage system contents to another appliance, which can help to facilitate disaster recovery. Data is transmitted securely via a private SSL connection to a remote Sun Storage 7000 Unified Storage System configured as remote replication target. Figure 10 shows a high availability configuration in which VMs can be replicated to remote storage, creating a disaster recovery scenario for both data and virtual machines.
- *Snapshots.* ZFS provides the Sun Storage 7000 Unified Storage System with an unlimited number of writable snapshots (subject to available capacity). Snapshots can be run manually or scheduled, and can even be exported as a new share.
- *Compression.* Storage system software provides four levels of compression, which is available to both NFS and CIFS filesystems.

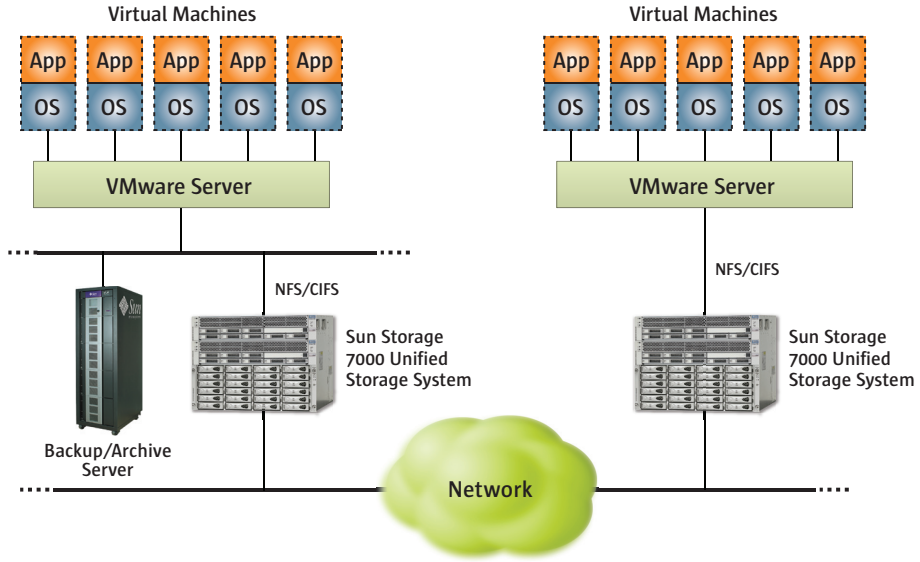

Figure 10. Sun Storage 7000 Unified Storage Systems enable robust data services, including remote replication

Core solution components

Table 2 details solutions for small, medium, and large configurations with a description of approximately the number of servers supported with each configuration. These sample solutions are designed to be comprehensive, but can be upgraded to meet site and business requirements. Additional servers and storage products fit seamlessly into these core reference architectures.

Table 2. Sun Storage 7000 Unified Storage System and VMware Solution Architectures for Windows Server Consolidation

	Small	**Medium**	**Large**
Description	Ideal for demonstrations or training staff new to VMware	Ideal for developers and unit testing	Ideal for architects as a staging environment (without full-blown production data)
Server Count	Approximately 10	Approximately 50	Approximately 100
Storage Configuration			
Connectivity	NFS	NFS	NFS
Storage Appliance	Sun Storage 7110 Unified Storage System	Sun Storage 7210 Unified Storage System	Sun Storage 7410 Unified Storage System (clustered for high availability if needed)
RAID Level Options	Striping, mirroring, and double parity	Striping, mirroring, and double parity	Striping, mirroring, and double parity
Capacity	2.0TB or 4.0TB	11.5TB to 142TB	24TB - 288TB (future 576TB); expandable
VMware Server			
Sun Server	Sun Fire X4170 server	Multiple Sun Fire X4170 servers	Sun Blade 6000 Server with 10 Sun X6270 processor blades, or multiple Sun Fire X4170 servers
Processor per system	Up to two Intel Xeon Processor 5500 Series	Up to two Intel Xeon Processor 5500 Series	Up to two Intel Xeon® Processor 5500 Series
Memory	8GB – 144GB	32GB – 144GB	72GB - 144GB
Networking	Four on-board Gigabit Ethernet ports	Four on-board Gigabit Ethernet ports, one dual, and quad port NIC controllers	2 1GbE and shared 2 10GbE NICs per blade; four on-board Gigabit Ethernet ports and two dual port NIC controllers in Sun Fire X4170 servers
Remote Management	Integrated Lights Out Management (ILOM)	Integrated Lights Out Management (ILOM)	Integrated Lights Out Management (ILOM)
Local Storage	2-16 146GB SAS disk	Two 146GB SAS HDDs	Two 146GB SAS HDDs

Disk Controller	Sun StorageTek SAS RAID HBA, Internal (SGXPCIESAS-R-INT-Z)	8-Port internal SAS PCI-Express (SG-PCIE8SAS-I-Z)	8-Port internal SAS PCI-Express (SG-PCIE8SAS-I-Z)
Internal RAID Level Options	RAID 0, 1, 1E, 5, 5EE, 6, 10, 50, 60	RAID 0 (Mirrored)	RAID 0 (Mirrored)
Local Capacity	VMware Vsphere 4 - 146GB (Mirrored) Virtual Machines (VM) ~ 2TB	VMware Vsphere 4 - 146GB (Mirrored)	VMware Vsphere 4 - 146GB (Mirrored)
Rack Mount	SESY9RK2Z	SESY9RK2Z	SESY9RK2Z
vCenter Server			
Server	Sun Fire X4170	Sun Fire X4170	Sun Fire X4170
Processor	Two Intel Xeon processors	Two Intel Xeon processors	Two Intel Xeon processors
Memory	24GB	24GB	24GB
Microsoft SQLServer			
Server	Sun Fire X4170	Sun Fire X4170	Sun Fire X4170
Processor	Two Intel Xeon processors	Two Intel Xeon processors	Two Intel Xeon processors
Memory	48GB	48GB	48GB
Disk Drives	2x 146GB	2x 146GB	2x 146GB
Backup Server (Optional)			
Tape Device		Sun StorageTek SL24	Sun StorageTek SL500
Software	N/A	NetBackup or other NDMP	NetBackup or other NDMP
VMware Software			
VMware server	VMware VI3 or VSphere 4.0	VMware VI3 or VSphere 4.0	VMware VI3 or VSphere 4.0
Hypervisor	VMware ESX and ESXi	VMware ESX and ESXi	VMware ESX and ESXi
Storage	VMDK	VMDK	VMDK
Multi-processor support for VMs	Virtual SMP	Virtual SMP	Virtual SMP
Management	vCenter License	vCenter License	vCenter License
Backup	Virtual Consolidated Backup (VCB), NetBackup, or other NDMP	Virtual Consolidated Backup (VCB), NetBackup, or other NDMP	Virtual Consolidated Backup (VCB), NetBackup, or other NDMP
Patch Management	VMware Update Manager	VMware Update Manager	VMware Update Manager
Availability	N/A	High Availability (HA)	High Availability (HA)
Live VM Migration	N/A	VMotion	VMotion
Live VM Disk Migration	N/A	Storage VMotion	Storage VMotion
Power Management	N/A	DPM	DPM
Resource Management	N/A	DRS	DRS

Custom configurations

Customers may need to customize the above solutions to further meet their needs. It is recommended that customers work with Sun-certified VMware experts to help modify the above configurations to fit specific needs.

A single point of contact for VMware solutions: software, storage, servers, and support

Sun combines best-of-class technologies with exceptional services and support for businesses that need servers, storage, and VMware software to reduce Microsoft server sprawl and consolidate servers. Sun now offers the VMware vSphere 4.0 product suite on Sun hardware systems, with full support from Sun available. Sun also provides a variety of software service plans that supply telephone and online technical support, software updates and upgrades, and self-help tools and notifications. A wide spectrum of service options are available with varying levels of support. Consultants from Sun Professional Services can also help to design and implement an end-to-end VMware infrastructure that is customized to meet virtualization needs.

In addition, Sun makes it easy to transition from an existing storage infrastructure to Sun Storage 7000 Unified Storage Systems. The Sun Unified Storage Data Migration Service provides the expertise needed to migrate storage systems safely, quickly, and securely. This service can help to overcome obstacles — such as lost time, administrative expenses, application downtime, and loss of productivity — that are often associated with data migrations.

For more information

To learn more about Sun Storage 7000 Unified Storage systems products and how Sun can streamline your Microsoft server and storage infrastructure, please contact a Sun sales representative. For more information, consult the related documents and Web sites listed in Table 3.

Table 3. Related Web sites

Web Site URL	Description
sun.com/vmwstorage	Sun storage and VMware
sun.com/vmware and vmware.com/sun	VMware and Sun partnership
www.sun.com/software/vmware/certification.jsp	Sun server and storage platforms certified for VMware infrastructure solutions
www.sun.com/solutions/virtualization/wsc	Sun solution for Windows consolidation
sun.com/unifiedstorage	Product information on the Sun Storage 7000 Unified Storage System
wikis.sun.com/display/BluePrints/An+Economic+Approach+to+Maximizing+Data+Availability	RAS functionality in the Sun Storage 7000 Unified Storage System
wikis.sun.com/display/BluePrints/Flash-Open+Storage+BluePrints	Technical best practices for open storage

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN (9786) Web sun.com