



Technical Report

Reference Architecture Design Guide

Microsoft Exchange Server, SQL Server, and SharePoint Server Mixed Workload on VMware vSphere 4, NetApp Unified IP Storage (NFS and iSCSI), and CISCO Nexus Unified Fabric

VMware and Microsoft Solutions Engineering, NetApp
August 2009 | TR-3785

TABLE OF CONTENTS

1	INTRODUCTION	4
2	SOLUTION SUMMARY	5
3	NFS SOLUTION DESIGN	6
3.1	HIGH LEVEL SOLUTION ARCHITECTURE	6
3.2	SOLUTION HARDWARE AND SOFTWARE REQUIREMENTS	6
3.2.1	HARDWARE RESOURCES	6
3.2.2	SOFTWARE RESOURCES	7
3.3	SOLUTION ARCHITECTURE DETAILS	8
3.3.1	VIRTUAL MACHINE LAYOUT	8
3.3.2	NETWORK ARCHITECTURE	8
3.3.3	STORAGE ARCHITECTURE	11
3.4	STORAGE SIZING	15
3.5	BACKUP AND RESTORE ARCHITECTURE	16
3.5.1	HIGH LEVEL BACKUP ARCHITECTURE	16
3.5.2	MICROSOFT SQL SERVER 2008	17
3.5.3	MICROSOFT EXCHANGE 2007	18
3.5.4	MICROSOFT OFFICE SHAREPOINT SERVER 2007	19
4	NFS REFERENCE ARCHITECTURE SOLUTION VALIDATION	19
4.1	STORAGE EFFICIENCY	19
4.2	PERFORMANCE VALIDATION	20
4.2.1	MICROSOFT EXCHANGE 2007	20
4.2.2	MICROSOFT SQL SERVER 2008	21
4.2.3	MICROSOFT SHAREPOINT SERVER 2007	21
4.3	VMWARE VMOTION, HA, AND DRS VALIDATION	21
4.4	BACKUP AND RESTORE VALIDATION	22
4.4.1	MICROSOFT EXCHANGE SERVER 2007	22
4.4.2	MICROSOFT SQL SERVER 2008	22
4.4.3	MICROSOFT SHAREPOINT 2007	22
5	ISCSI REFERENCE ARCHITECTURE DESIGN	23
5.1	HIGH LEVEL SOLUTION ARCHITECTURE	23
5.2	SOLUTION HARDWARE AND SOFTWARE REQUIREMENTS	23
5.2.1	HARDWARE RESOURCES	23
5.2.2	SOFTWARE RESOURCES	24
5.3	ISCSI SOLUTION ARCHITECTURE DETAILS	25
5.3.1	VIRTUAL MACHINE LAYOUT	25

5.3.2	NETWORK ARCHITECTURE	25
5.3.3	STORAGE ARCHITECTURE	28
5.3.4	STORAGE SIZING	29
5.4	BACKUP AND RESTORE ARCHITECTURE	30
5.4.1	BACKUP ARCHITECTURE DETAILS	30
5.4.2	RESTORE ARCHITECTURE DETAILS	30
6	ISCSI REFERENCE ARCHITECTURE SOLUTION VALIDATION	31
6.1	STORAGE EFFICIENCY	31
6.2	PERFORMANCE VALIDATION	31
6.2.1	MICROSOFT EXCHANGE 2007	32
6.2.2	MICROSOFT SQL SERVER 2008	32
6.2.3	MICROSOFT SHAREPOINT SERVER 2007	32
6.3	VMWARE VMOTION, HA, AND DRS VALIDATION	33
6.4	BACKUP AND RESTORE VALIDATION.....	33
6.4.1	MICROSFT EXCHANGE SERVER 2007	33
6.4.2	MICROSOFT SQL SERVER 2008	33
6.4.3	MICROSOFT OFFICE SHAREPOINT SERVER 2007	34
7	SUMMARY	34
8	ACKNOWLEDGEMENTS	34
9	FEEDBACK.....	35

1 INTRODUCTION

As customers move toward their goal of 100% virtualized data centers, they increasingly look for ways to bring the benefits of VMware® virtualization to their mission-critical Microsoft® applications. Customers who are planning a new deployment, performing an upgrade, or planning to 100% virtualize their data centers have an ideal opportunity to transition to a [VMware vSphere™](#) virtual infrastructure built on [NetApp® unified storage](#).

This document provides guidance on how to successfully design and architect a scalable Microsoft applications mixed workload solution on highly available VMware vSphere 4 virtual infrastructure and NetApp unified storage. It highlights leveraging the IP storage-based solutions, iSCSI and NFS for hosting virtual machines. It also describes the NetApp backup and recovery solution for the Microsoft applications. Both the iSCSI- and NFS-based solutions are applicable for all enterprise types (large, midsize, and SMB) and can be scaled up or down based on business requirements. Some key benefits of the overall solution are:

- **Reduced costs with VMware virtualization.** For many organizations, upgrading to newer Microsoft server applications without virtualization can result in investing more server hardware in an application that has already become excessively costly to run. VMware virtualization can unlock the full power of the hardware by running multiple workloads on each system. This can provide a cost-effective solution and quicker ROI when compared to deployments without virtualization.
- **Advanced NetApp unified and efficient storage solutions.** Customers can deploy Microsoft Exchange, SQL Server®, and SharePoint® on storage solutions that leverage existing networking infrastructure such as iSCSI and NFS, which can offer a very cost-effective storage solution. NetApp FAS and V-Series storage arrays have been fully tested and certified for use in IP-based (both iSCSI and NFS) VMware environments. Also, by leveraging NetApp storage efficiency and intelligent caching capabilities, customers can save significantly on their storage investment without negative tradeoffs.
- **High availability.** A VMware-enabled platform can provide high availability (HA) for Microsoft server applications without the need for clustering at the virtual machine (VM) level. Virtual machines are no longer tied to the underlying server hardware and can be moved across servers at any time with VMware VMotion®. VMware HA provides server hardware fault tolerance for every VM and offers greater levels of availability over solutions designed to protect just the server.
- **Advanced backup and recovery solutions.** The NetApp backup and recovery solution is built using integrated VMware, Microsoft, and NetApp technologies for advanced, application-aware data protection. Deduplication-aware remote replication for disaster recovery with NetApp SnapMirror® provides an end-to-end data protection solution.

The primary Microsoft applications virtualized are:

- [Microsoft Exchange Server 2007](#)
- [Microsoft Office SharePoint Server 2007](#)
- [Microsoft SQL Server 2008](#)

The key highlights of this architecture are:

- Microsoft applications virtualization with [VMware vSphere 4](#)
- Storage efficiency with NetApp primary storage [deduplication](#), and thin provisioning, without any negative tradeoffs
- Scalability and ease of management with NetApp deduplication-aware iSCSI and NFS datastores, VMware vSphere 4, NetApp unified storage, and backup and recovery solution
- Efficient, deduplication-aware, application-consistent backup and recovery with NetApp [SnapManager® for Virtual Infrastructure \(SMVI\)](#), [SnapManager for Exchange \(SME\)](#), [SnapManager for SQL Server \(SMSQL\)](#), [SnapManager for Microsoft Office SharePoint Server \(SMOSS\)](#), and NetApp [SnapMirror](#) replication for iSCSI-based solution
- Efficient, deduplication-aware, application-consistent backup and recovery with VMware snapshots leveraging VMware VSS, NetApp SMVI, and NetApp SnapMirror replication for NFS-based solution

For more information about the best practices followed in this architecture, please see the following guides:

- [TR-3749](#): NetApp and VMware vSphere 4 Storage Best Practices
- [TR-3730](#): SnapManager 5.0 for Microsoft Exchange Best Practices Guide
- [TR-3578](#): Microsoft Exchange Server 2007 Best Practices Guide
- [TR-3768](#): SnapManager 5.0 for SQL Server: Best Practices Guide
- [TR-3715](#): SnapManager for Microsoft Office SharePoint Server: Backup and Recovery Guide
- [TR-3747](#): NetApp Best Practices for File System Alignment in Virtual Environments
- [TR-3505](#): NetApp Deduplication for FAS and V-Series Best Practices Guide
- [TR-3737](#): SMVI 1.0 Best Practices

2 SOLUTION SUMMARY

The solution showcases virtualizing Microsoft applications on VMware vSphere 4 virtual infrastructure and NetApp unified storage, achieving the desired level of storage efficiency, performance, operational agility, and efficient, application-consistent, deduplication-aware data protection. Results of the testing demonstrate that the performance of Microsoft applications on VMware vSphere and NetApp storage in this solution is suitable for production environments and is well within Microsoft best practice recommendations.

Features of the VMware vSphere platform, including VMware VMotion, HA, and Distributed Resource Scheduler (DRS), were tested and demonstrated substantial increases in overall flexibility and availability. NetApp SMVI, along with VMware VSS-aware snapshots, was used as the solution for backup and restore of the virtualized servers in the NFS environment, and SMVI, SME, SMSQL, and SMOSS were used as the solution for application-consistent backup and restore of the virtualized servers in the iSCSI environment. NetApp SnapMirror was used for replication of the backup data to a secondary NetApp storage controller to provide disaster recovery, showing the flexibility of backup and recovery options available to customers with the NetApp solution. This guide describes the solution architecture and includes hands-on test results of typical administrative activities such as application backup and restore.

The workload virtualized in both the NFS and iSCSI architectures is as follows:

- **Microsoft Exchange 2007**. 3,000 heavy users with 250MB mailbox per user, and 0.33 IOPS per user
- **MS SharePoint 2007**, 3,000 users with 390MB space per user
- **SQL Server 2008**. 3,000 users, 10 databases using OLTP, DSS, and mixed workloads such as CRM (sales and manufacturing) etc.

Table 1 shows the virtual CPU and memory configuration for Exchange, SharePoint, and SQL Server 2008 VMs.

Table 1) Exchange, SQL Server, and SharePoint VM configuration.

Microsoft Application	Virtual Machine	Virtual CPU	Memory (GB)
Exchange 2007	(2) Exchange Mailbox Servers	2	9.5
	(2) Exchange CAS Servers, (2) HUB Servers	1	1
SQL Server 2008	(2) SQL Server	4	4
SharePoint Server 2007	(2) Web Servers	2	2
	(1) Index Server, (1) SQL Server	4	4

The solution also includes Microsoft IIS Web servers, and test and development servers, along with the following solution-monitoring and management tools from NetApp, VMware and Microsoft.

- NetApp DataFabric[®] Manager (Operations Manager)
- VMware vCenter[™] 4
- Microsoft System Center Operations Manager (SCOM), and Windows Server[®] Update Services Server (WSUS)

Both the NFS and iSCSI based solutions are scalable and can be easily tailored for deployment of any size.

Note: Hardware requirements vary from environment to environment, based on specific workload requirements.

3 NFS SOLUTION DESIGN

3.1 HIGH LEVEL SOLUTION ARCHITECTURE

As shown in Figure 1, the infrastructure used for this solution validation involved three VMware ESX 4.0 hosts running the mixed Microsoft applications workload described in section 2, with a total of 30 virtual machines hosted on NetApp shared storage. The VM operating system, installed applications, databases, and logs are all hosted on NetApp-based NFS datastores. Solution high availability is achieved by using VMware HA, NetApp active-active controllers, and Cisco Nexus 5020 switches. The backup and recovery solution component includes application-consistent point-in-time NetApp Snapshot™ copies with NetApp SMVI and SnapMirror replication to a secondary site.

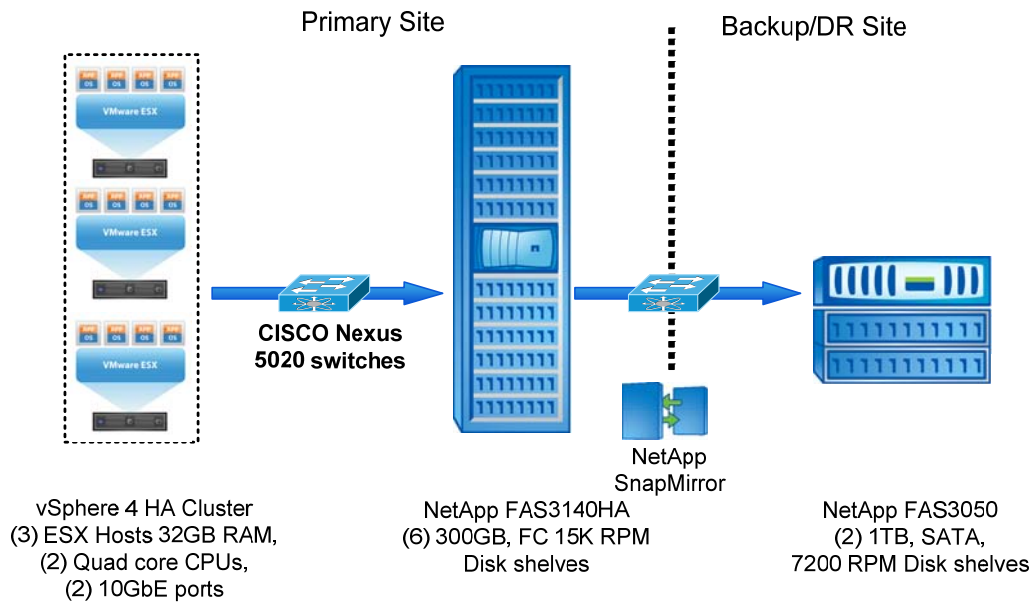


Figure 1) High-level solution architecture

3.2 SOLUTION HARDWARE AND SOFTWARE REQUIREMENTS

3.2.1 HARDWARE RESOURCES

The following equipment was used in this configuration validation.

Table 2) Hardware configuration

Solution Component	Minimum Revision
Primary Storage	
(1) NetApp FAS3140HA Cluster	Data ONTAP® 7.3.1.1
(6) disk shelves	84 disks (14 per shelf); each disk 300GB/15K/FC Note: To validate the NFS solution, six shelves were available, as compared to five shelves for the iSCSI solution. Therefore there are some differences in how aggregates were used for hosting the VM OS, vswap, and page files. This should not be considered limitation of the protocol.

Solution Component	Minimum Revision
Networking	
(2) Cisco Nexus 5020 switches	
(1) dual port 10Gb Ethernet NIC per FAS3140 controller	
Backup Storage	
(1) NetApp FAS3050HA cluster	Data ONTAP® 7.3.1.1
(2) Disk shelves	28 disks (14 per shelf); each disk 1TB/7200RPM/SATA
(3) ESX hosts	
32GB RAM	
(2) Quad Core Xeon™ Processors	
(1) dual port 10Gb Ethernet NIC	

3.2.2 SOFTWARE RESOURCES

The following software components were used in the configuration validation.

Table 3) Software configuration

Solution Component	Minimum Revision
Primary Storage	
Data ONTAP	7.3.1.1
NFS, ASIS, FlexClone®, SnapMirror®, SnapRestore®, and NearStore® Licenses	N/A
Backup Storage	
Data ONTAP	7.3.1.1
NFS, ASIS, SnapMirror®, and NearStore Licenses	N/A
NetApp Management Software	
NetApp Virtual Service Console (VSC)	1.0
NetApp System Manager	1.0
NetApp SnapManager for Virtual Infrastructure	1.2
NetApp DataFabric Manager	3.8
VMware vSphere Infrastructure	
ESX Hosts	VMware ESX , 4.0.0 (build 164009)
vCenter Server	4.0.0
vCenter Database	SQL Server 2005

Solution Component	Minimum Revision
Applications Virtual Machine Operating System	
Windows Server 2003 Note: Current implementation of the VSS requestor in VMware Tools does not support Windows Server 2008 from the perspective of application consistency. For more information, see the Virtual Machine Backup Guide from VMware.	x64, Enterprise Edition, Service Pack2
Microsoft Applications	
Microsoft Exchange Server	2007, Enterprise Edition, SP1
Microsoft Office SharePoint Server	2007 Enterprise Edition SP2
Microsoft SQL Server	2008 Enterprise Edition

3.3 SOLUTION ARCHITECTURE DETAILS

3.3.1 VIRTUAL MACHINE LAYOUT

The solution described in this section uses a total of 30 virtual machines. This configuration simulates a real-world customer environment with the supporting utility and test and dev servers in addition to the primary Microsoft application servers.

MICROSOFT APPLICATIONS VMS

- **Microsoft Exchange Server 2007.** Total 6 VMs (2 Mailbox servers, 2 HUB servers, 2 CAS servers)
- **Microsoft Office SharePoint Server 2007.** Total 4 VMs (2 Web front end, 1 search/index, 1 SQL Server 2008)
- **Microsoft SQL Server 2008.** 2 VMs
- **Microsoft IIS.** 4 VMs

TEST AND DEV VMS

- **Windows Server 2003.** 4 VMs

UTILITY VMS

(1) Microsoft WSUS, (1) Microsoft SCOM, (1) Microsoft Exchange LoadGen Tool, (4) SharePoint Test Workstations, (1) NetApp DataFabric Manager, (1) VMware vCenter 4, (1) VMware vCenter 4 database

All of the virtual machines were aligned using the procedures described in NetApp [TR3747](#): "Best Practices for File System Alignment in Virtual Environments." Also, all of the NetApp recommended NFS settings described in NetApp [TR3749](#), "NetApp vSphere Best Practices," were set using the NetApp VSC vCenter plugin, directly from the vCenter GUI.

3.3.2 NETWORK ARCHITECTURE

In this solution, the network comprised of two Cisco Nexus 5020 switches. Since the Cisco Nexus switches used in this configuration support virtual port channeling (vPC), logical separation of the storage network from the rest of the network is achieved while at the same time providing a high level of redundancy, fault tolerance, and security. With the vPC feature, scalable Layer 2 topologies can be deployed, reducing the dependence on Spanning Tree Protocol for redundancy and loop avoidance. Also, high cross-sectional bandwidth is attained by the features ability to use all available physical links that interconnect the devices.

On the Nexus network ensure the following configurations.

- Be sure to setup a management VLAN for the service console, a public VLAN for the virtual machine network, and a private, non-routable VLAN for VMotion™
- Be sure to use a 10Gb connection between the two Nexus 5020 switches
- Be sure to enable a vPC between the two Nexus 5020 switches. In order to use this feature please be sure to have the Cisco NX-OS Software Release 4.1(3)N1 for Cisco Nexus 5000 series switches installed on your Cisco Nexus 5020

While the Cisco Nexus 5020 switches are 10Gb, they do support 1Gb modules. Therefore, other Cisco switches can be used in conjunction with the Nexus 5020's in order to further scale out a virtualization and storage network.

ESX HOST NETWORK ARCHITECTURE

Figure 2 shows the virtual network layout for each ESX host. Each ESX host has two 10 Gb Ethernet ports configured into different port groups as shown in the figure.

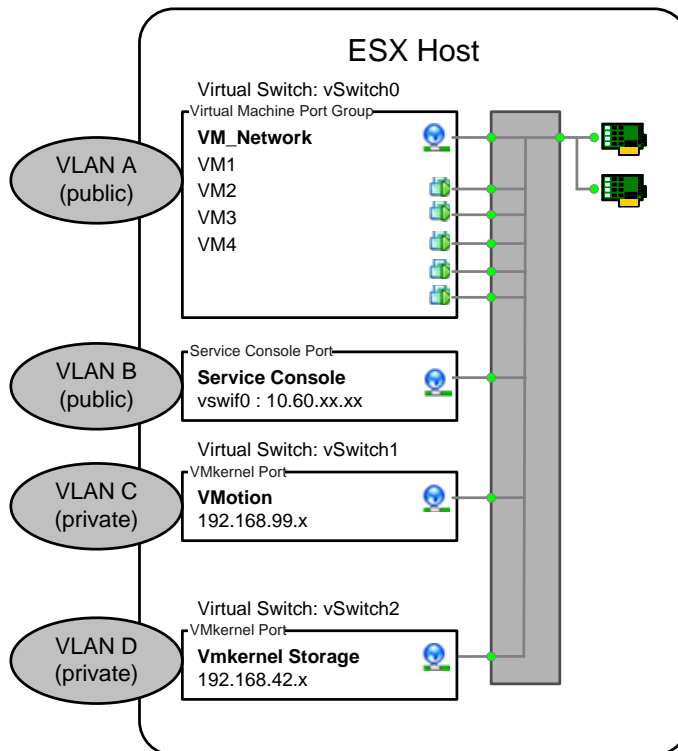


Figure 2) ESX Host network architecture

STORAGE NETWORK LAYOUT

Figure 3 shows the storage network layout for the ESX host connectivity with the NetApp storage controller over Cisco Nexus 5020 switches. Make sure to configure a non-routable VLAN for the NFS storage traffic to pass to and from the NetApp storage controllers to the vSphere hosts. With this setup, the NFS traffic is kept completely contained, and security is more tightly controlled.

Also, it is important to have at least two physical Ethernet switches for proper network redundancy in your VMware environment.

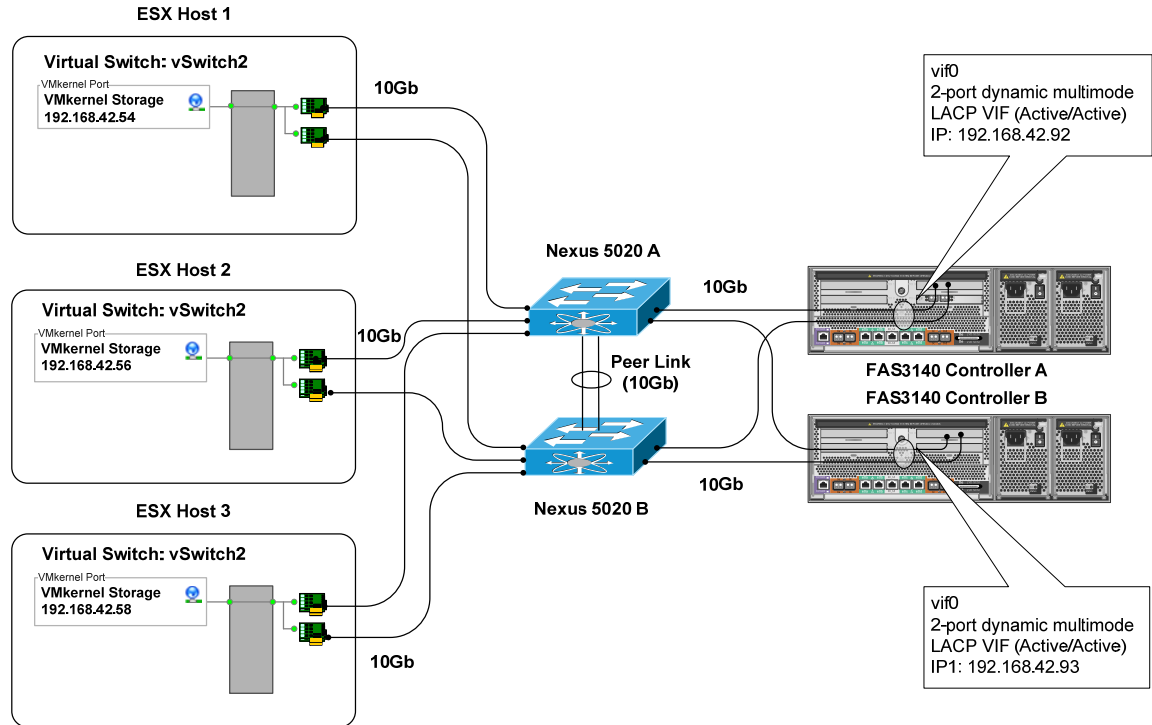


Figure 3) Storage network architecture.

3.3.3 STORAGE ARCHITECTURE

NETAPP STORAGE AGGREGATE LAYOUT

Figure 4 shows the NetApp storage aggregate layout for hosting the different data components for every VM. NetApp aggregates provide a large virtualized pool of storage capacity and disk IOPS to be used on demand by all the virtual machines hosted in the aggregate. This can be compared to the VMware virtualization where CPU and memory resources are pooled and leveraged on demand.

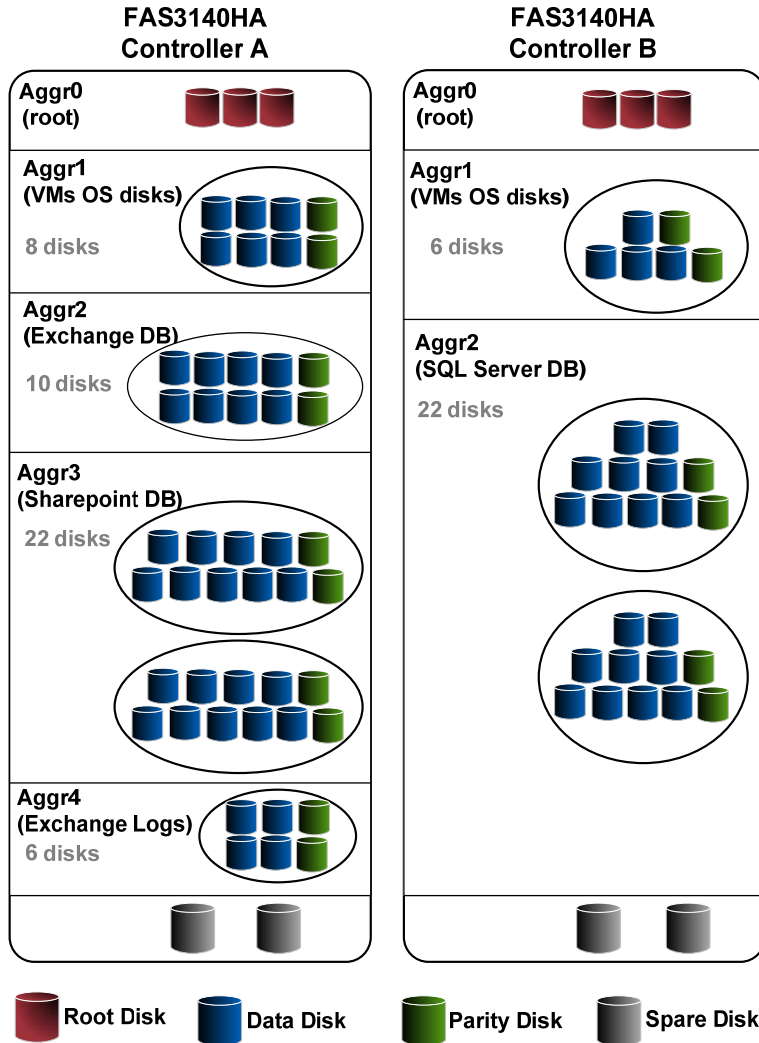


Figure 4) NetApp storage aggregate layout.

The aggregate sizing is based on the storage requirements for all the applications to meet the storage capacity, performance, and snapshot backup requirement of an assumed workload. When sizing for your environment, please consult with your NetApp SE about the exact storage configuration based on your individual requirements.

Note: In this solution, all of the aggregates hosting volumes required for SharePoint and Exchange VMs are hosted on one storage controller and the aggregates hosting volumes for SQL Server are hosted on the second controller. This consideration was made from the perspective of the [VMware vCenter Site Recovery Manager](#), which we plan to add in the future release of this guide. The VMware vCenter Site Recovery Manager requires all datastores hosting data for a VM to be on the same storage controller.

Separate aggregates have been created on each controller for hosting the OS data for different VMs (VM C drives, page file, and VM vswap file). Although this may not be necessary for a small configuration with few VMs, it should be considered when the environment will scale to hundreds to thousands of VMs.

Also note that separate aggregates have been created for each application workload. Although this is not required and multiple application workloads can be served off the same aggregate, the reason here is to show the storage design for each application in a modular way. This will be helpful if you plan to virtualize one of these applications in your environment.

NETAPP STORAGE VOLUME LAYOUT

Figure 5 shows the NetApp storage volume layout for hosting the different data components for every VM.

Note: All of the volumes were thin provisioned to use the capacity on demand.

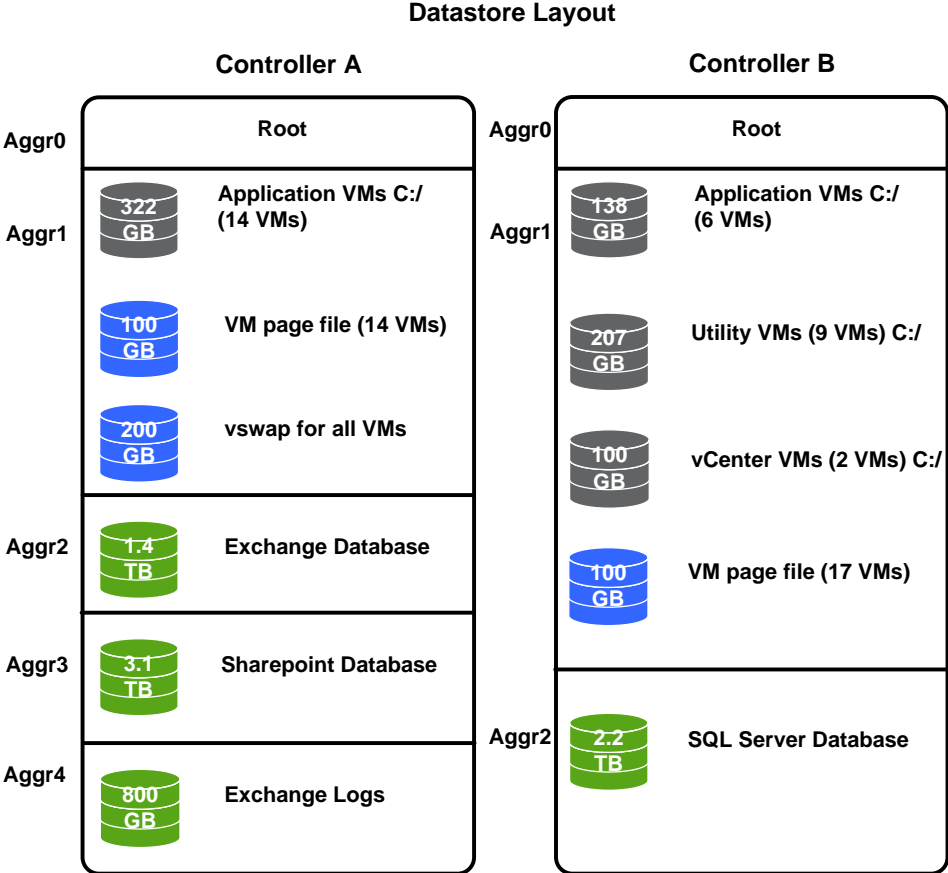


Figure 5) NetApp storage volume layout

MICROSOFT EXCHANGE SERVER 2007 DATASTORE LAYOUT

Figure 6 shows the datastore layout for the different data components of Microsoft Exchange Server 2007. The temporary data – VM page file and VM swap file (.vswp) – have been separated out. This will reduce the daily snapshot change rate and facilitate faster completion of nightly primary storage deduplication operations. The database and transaction logs vmdk files are hosted on separate datastores on separate aggregates, as shown in Figure 5.

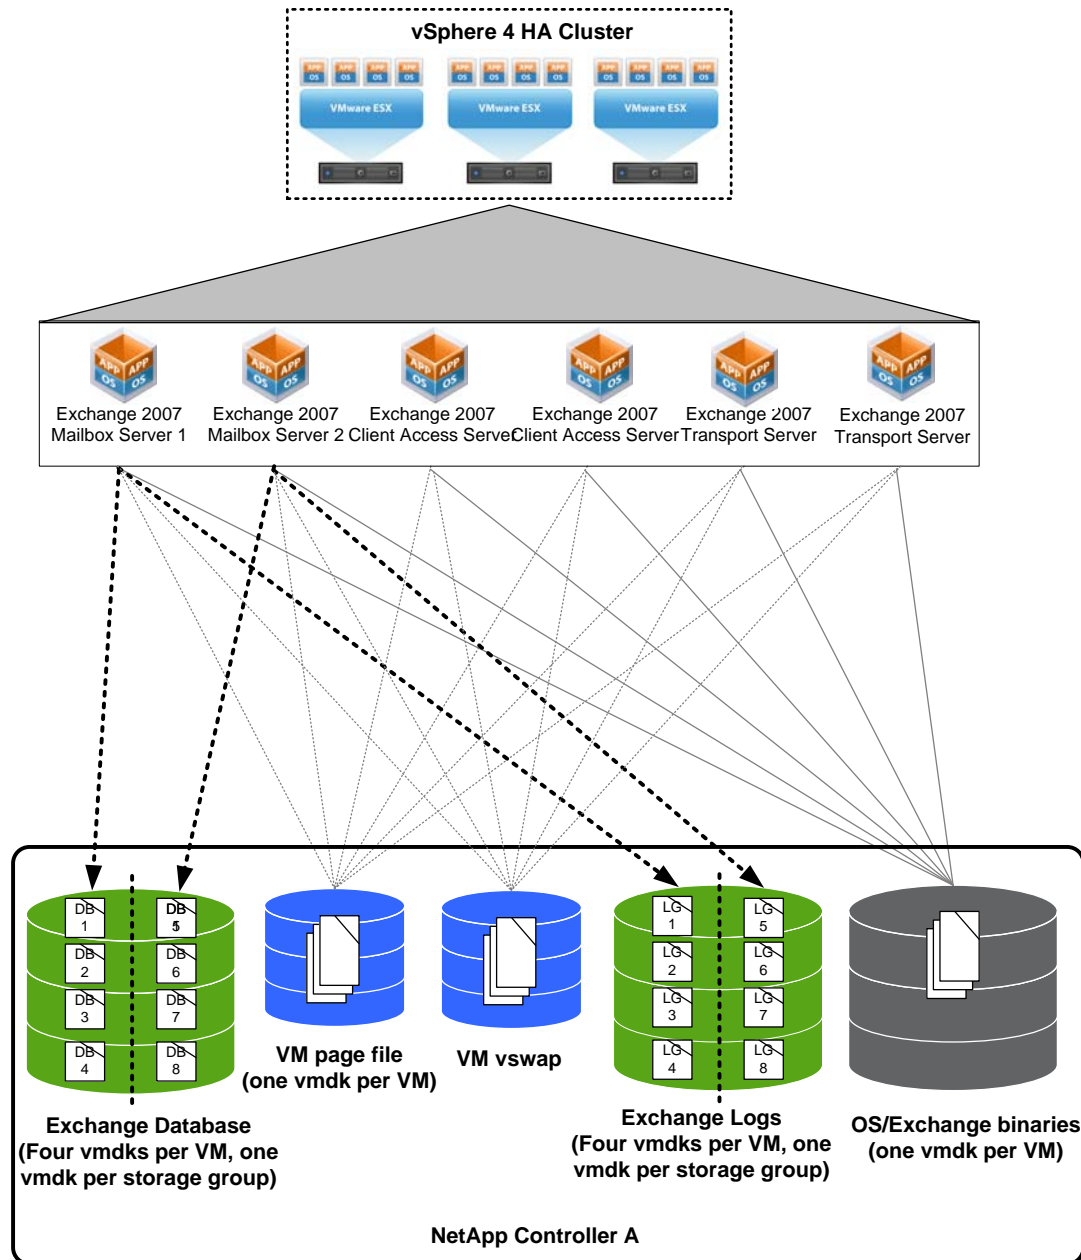


Figure 6) Microsoft Exchange Server 2007 datastore layout

MICROSOFT OFFICE SHAREPOINT SERVER 2007 DATASTORE LAYOUT

Figure 7 shows the datastore layout for the different data components of Microsoft Office SharePoint Server 2007. The database and log files are hosted on a separate datastore than the Windows® OS and SharePoint binaries. This datastore is hosted on a dedicated aggregate, as shown in Figure 5.

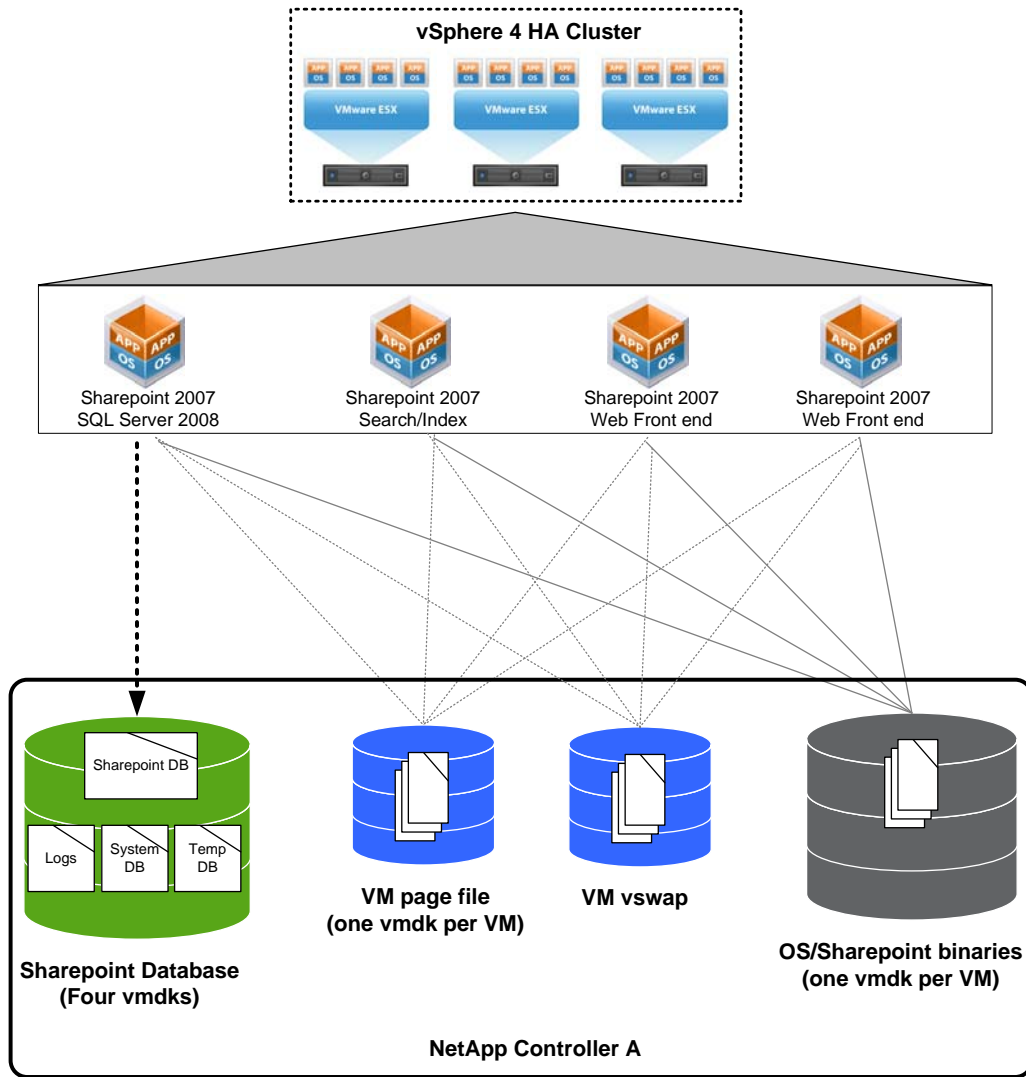


Figure 7) Microsoft Office SharePoint Server 2007 datastore layout

MICROSOFT SQL SERVER 2008 DATASTORE LAYOUT

Figure 8 shows the datastore layout for the different data components of Microsoft SQL Server 2008. The SQL database and log files are hosted on a separate datastore than the Windows OS and SQL Server binaries.

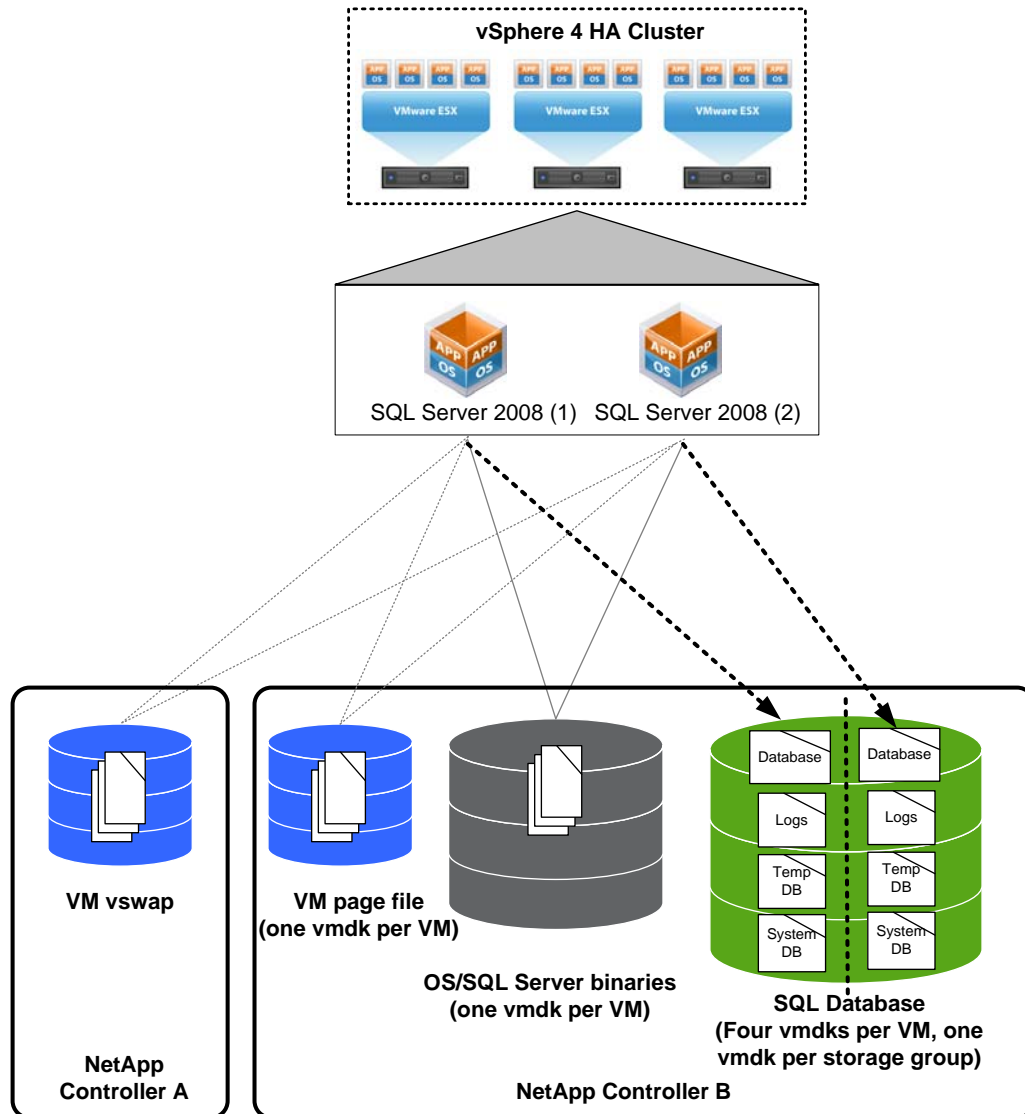


Figure 8) Microsoft SQL Server 2008 datastore layout

3.4 STORAGE SIZING

This section describes the storage sizing. These numbers vary from environment to environment, and you should consult your NetApp systems engineer about the exact sizing for your environment.

- **Microsoft SQL Server 2008.** The SQL Server workload was divided into 10 separate databases using both OLTP and DSS workloads. The databases required 2.2TB of disk space and the transaction logs required 4.5GB of disk space.
- **Microsoft Office SharePoint Server 2007.** SharePoint used 3TB of disk space for the database and 1GB of disk space for the transaction logs.

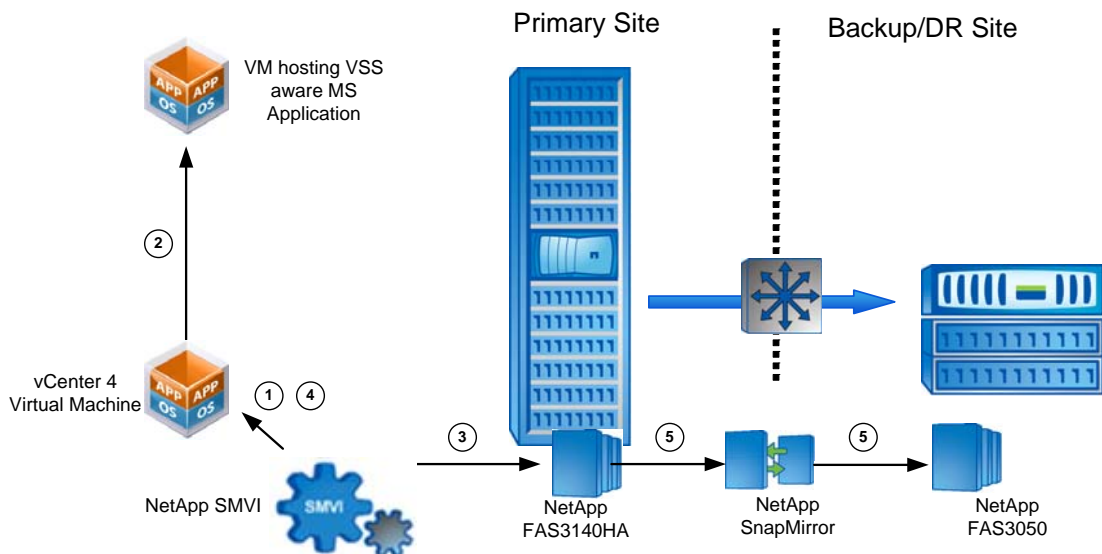
- **Microsoft Exchange 2007.** Exchange used 1TB of disk space for the databases and 240GB of disk space for the transaction logs.
- **Datastores hosting VM C: drives (OS and application binaries).** (Size of the VM C: drive + 15% free space for VMware snapshots and non-vmdk files) * Number of VMs in the datastore
 - Application VM datastore on Controller A, hosting 14 VMs = 322GB
 - Application VM datastore on Controller B, hosting 6 VMs = 138GB
 - Utility VM datastore on Controller B, hosting 9 VMs = 207GB

Note: These storage requirements are before considering NetApp deduplication and thin provisioning. The NetApp deduplication and storage efficiency potential savings are discussed in section 4. The deduplication savings vary from environment to environment and application to application. For more information, see NetApp [TR3505](#).

3.5 BACKUP AND RESTORE ARCHITECTURE

3.5.1 HIGH LEVEL BACKUP ARCHITECTURE

Figure 9 shows the backup process workflow. NetApp SMVI is used as a backup and restore solution for the environment. Remote replication is achieved using NetApp SnapMirror, which is tightly integrated with NetApp SMVI. This makes it possible to invoke local snapshot backup and remote replication from a single user interface, using a single process. For SMVI best practices, see NetApp [TR3737](#). SMVI leverages the VSS requestor in VMware Tools to create application-consistent backups. This is invoked as part of the VMware snapshot performed before taking the Snapshot copy on the NetApp array.



SMVI Backup Workflow

1. SMVI invokes VMware snapshot for the VM hosting MS Application
2. VMware snapshot invokes VSS snapshot on the VM
3. SMVI invokes NetApp snapshot on the primary storage
4. SMVI invokes removing VMware snapshot for the VM
5. SMVI invokes SnapMirror update to the remote storage controller

Figure 9) High level backup architecture

Note: The current VSS requestor implementation in VMware Tools does not invoke transaction log truncation for the SQL database after backups. This could result in the log full/disk out of space condition as the transaction logs keep growing. However, transaction logs are required for achieving individual transaction-level recovery.

An automated way of achieving transaction log truncation is to change the database recovery model from Full to Simple. However, this means that only point-in-time database recovery is possible. Individual transaction-level recovery cannot be achieved with the simple recovery model. Therefore this can be an excellent backup and recovery solution for databases that do not require individual transaction log level recovery.

For more information about managing the disk space for SQL transaction logs, see the following VMware KB article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1001046

3.5.2 MICROSOFT SQL SERVER 2008

NetApp SMVI was used as a backup solution for achieving file system and application-consistent backups for both the OS and application data. This solution allows full VM and individual database-level recovery.

BACKUP ARCHITECTURE DETAILS

One backup policy is used to back up all the desired datastores associated with the SQL Server 2008 VMs. The backup application consistency is achieved by leveraging the VSS requestor in VMware Tools as part of the VMware snapshot that is initiated by NetApp SMVI before taking NetApp array based block-level Snapshot copies.

Achieving Transaction-Level Recovery

For databases where individual transaction-level recovery is required and transaction log truncation is automatically performed, NetApp highly recommends using NetApp SMSQL as your backup and recovery solution. Currently SMSQL is supported on iSCSI- and FC-based solutions. For details on using SMSQL as the backup and recovery solution for achieving transaction-level recovery, see section 6. Until the SMSQL support for NFS becomes available, when transaction-level recovery is required, you can leverage traditional file-based backup solutions from NetApp partners such as Symantec, CommVault, and IBM by installing a SQL agent inside the guest VM.

RESTORE ARCHITECTURE DETAILS

Restores from Local Snapshot Backups

NetApp SMVI allows full VM and individual database-level recovery. Full VM-level recovery (OS, application, and database) is achieved by using NetApp SMVI restore functionality. For individual, point-in-time database-level recovery, follow this process:

1. Leverage NetApp SMVI to mount the backup datastore to the ESX host running the SQL Server VM.
2. Attach the vmdk files hosting SQL database files (.mdf, .ndf, .ldf) to the SQL Server VM as new drives.
3. Detach the database that will be restored from backup by using SQL Server Management Studio.
4. Copy the .mdf, .ndf, and .ldf files for the database from the mounted drive to the original drives, provisioned for hosting the database files and transaction logs.
5. Attach the database by using SQL Server Management Studio. Be sure to specify the correct paths for the .mdf, .ndf, and .ldf files that were restored.

Restores from Remote SnapMirror Backups

For a full VM-level restore from the SnapMirror backups on the remote site, follow this process:

1. Quiesce and break the SnapMirror relationships
2. Set up a SnapMirror relationship back from the destination storage system on the remote site to the source storage system at the primary site.
3. Quiesce and break the new SnapMirror relationship again
4. Mount the datastores on the ESX host.
5. Register the VMs from the restored datastore.

Note: Make sure that all the VM hard drives point to the correct vmdk files on the restored datastores.

3.5.3 MICROSOFT EXCHANGE 2007

BACKUP ARCHITECTURE DETAILS

In the NFS architecture, NetApp SMVI was used as a backup solution for achieving both file system and application-consistent backups. Full VM, individual database, and single and multiple mailbox-level recovery can be achieved using this process. One backup policy is used to back up all the desired datastores associated with the Exchange Server 2007 VMs. Exchange Server application consistency is achieved by leveraging the VSS requestor in VMware Tools as part of the VMware snapshot initiated by NetApp SMVI before taking NetApp array-based block-level Snapshot copies.

Note: The current VSS requestor implementation in VMware Tools does a VSS copy backup instead of a VSS full backup. This implies that transaction logs don't get truncated after backups, which could result in a disks out of space condition as the transaction logs keep growing. Also, no automated backup verification can be performed using this process.

Alternate Backup and Recovery Solutions

NetApp SME along with single mailbox recovery can be leveraged to achieve application-consistent backups, enable transaction log truncation, and perform automated and granular recovery (up to single mailbox level recovery). At this point, SME is supported on iSCSI and FC protocol based datastores. Section 6 provides details on using SME in your environment. Until the SME support for NFS becomes available, here are the options for achieving transaction log truncation for Exchange Mailbox VMs hosted on NFS datastores

- In order to truncate the logs, a full backup of the Exchange Server needs to be done. This can be performed by using the NTBackup application provided in Windows Server 2003. Using NTBackup, a full backup of Exchange mailbox storage groups and databases can be performed; and the associated logs get truncated. If multiple Exchange mailbox storage groups and databases exist, better performance and decreased backup times can be achieved by creating a separate job for each database that run simultaneously. Be sure to create a separate VMDK disk drive on shared storage that will be used to host the backup files that are created. Also be sure to set the backup to Replace instead of Append. Figure 10 shows the workflow on how this should be used. For more information on using NTBackup for full backups of Exchange Server 2007, go to <http://technet.microsoft.com/en-us/library/aa998870.aspx>.

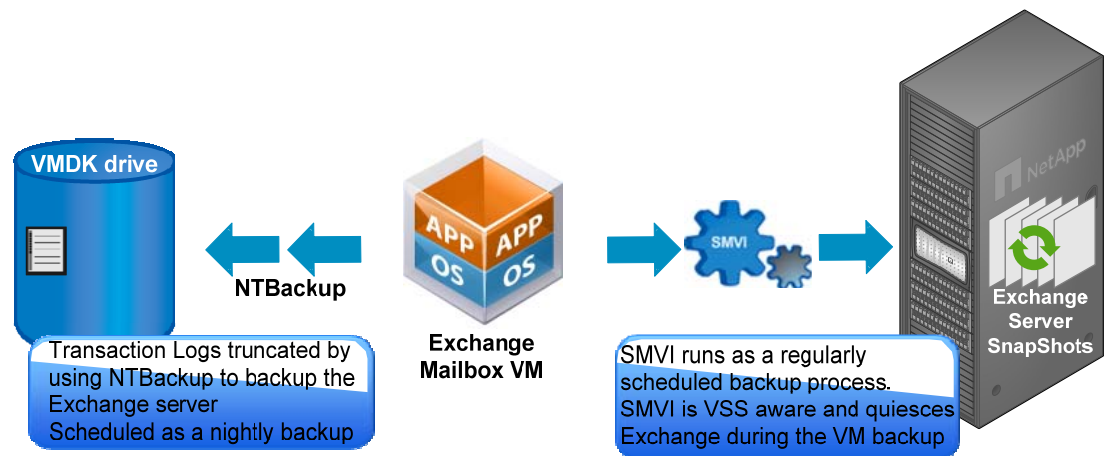


Figure 10) Leveraging NTBackup for transaction log truncation.

To validate that a successful backup of Exchange Server 2007 has occurred, set SMVI to send an e-mail upon completion of the job. Also, when using NTBackup to perform full backups for log truncation, review the Windows Server 2003 logs. Finally, to manually verify backups, run `eseutil /k` against the DB file and the logs. For information on `eseutil`, see <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=a0565435-6e79-41b0-a969-5766536db412>.

- Leverage traditional file-based backup solutions from NetApp partners, such as Symantec, CommVault, and IBM, by installing an Exchange agent inside the guest VM.

RESTORE ARCHITECTURE DETAILS

NetApp SMVI-based backup solution supports individual or multiple VM level restores. Individual mailbox restores can be performed using the built in Exchange Individual Mailbox restore tools.

3.5.4 MICROSOFT OFFICE SHAREPOINT SERVER 2007

BACKUP ARCHITECTURE DETAILS

NetApp SMVI is used as a backup solution for achieving both file system and application-consistent backups. One backup policy is used to back up all the VMs and desired datastores associated with all the VMs in the SharePoint farm. The SharePoint SQL Server application consistency is achieved by leveraging the VSS requester in VMware Tools as part of the VMware snapshot initiated by NetApp SMVI before taking NetApp array-based block-level Snapshot copies.

RESTORE ARCHITECTURE DETAILS

Restores from Local Snapshot Backups

The NetApp SMVI-based backup solution supports environment-level point-in-time recovery. Using SMVI restore functionality, all the VMs can be restored from the same backup. For more granular recovery, such as item level, site level, and so on, NetApp highly recommends leveraging the NetApp SMOSS. SMOSS leverages NetApp Snapshot copies and at the same time can provide granular recovery. SMOSS is currently supported on iSCSI and FC based solutions. Section 6 provides details on using SMOSS as the backup and recovery solution for achieving granular item level recovery.

Until SMOSS support for NFS becomes available, and more granular recovery is required for SharePoint environment hosted on NFS datastores, NetApp recommends leveraging traditional file-based backup solutions from NetApp partners, such as Symantec and CommVault by installing SharePoint specific agents inside the guest VM.

Restores from Remote SnapMirror Backups

Full environment-level restores from the SnapMirror backups can be achieved as follows:

- 1 Quiescing and breaking the SnapMirror relationships for all the volumes hosting the SharePoint data
- 2 Setting up the SnapMirror relationship from the destination storage system on the remote site back to the source storage system on the primary site
- 3 Quiescing and breaking this SnapMirror relationship (from step 2)
- 4 Mounting the datastores on the ESX host, and registering the VMs from the restored datastore

Note: Make sure that all the VM hard drives point to the correct vmdk files on the restored datastores.

4 NFS REFERENCE ARCHITECTURE SOLUTION VALIDATION

4.1 STORAGE EFFICIENCY

NetApp thin provisioning capabilities were enabled on all the datastores on the primary storage, and deduplication was enabled on the datastores hosting the VM OS and application binaries. The deduplication schedule was set to run once every night. Figure 11 shows the [NetApp System Manager](#) that was used to manage the storage, showcasing 92% savings for a datastore hosting OS and application binaries for 14 VMs hosting 6 Exchange VMs, 4 SharePoint VMs, and 4 test and dev VMs. Similar storage savings were observed for other datastores on Controller B, hosting the OS and application binaries for the SQL Server, IIS, utility, and vCenter VMs.

As you scale out to virtualize your entire data center with hundreds to thousands of VMs, the storage efficiency can even be higher. Also note that NetApp's intelligent caching capabilities (built natively in Data ONTAP and [Performance Accelerator Module](#) cards) strongly complement NetApp' storage efficiency capabilities.

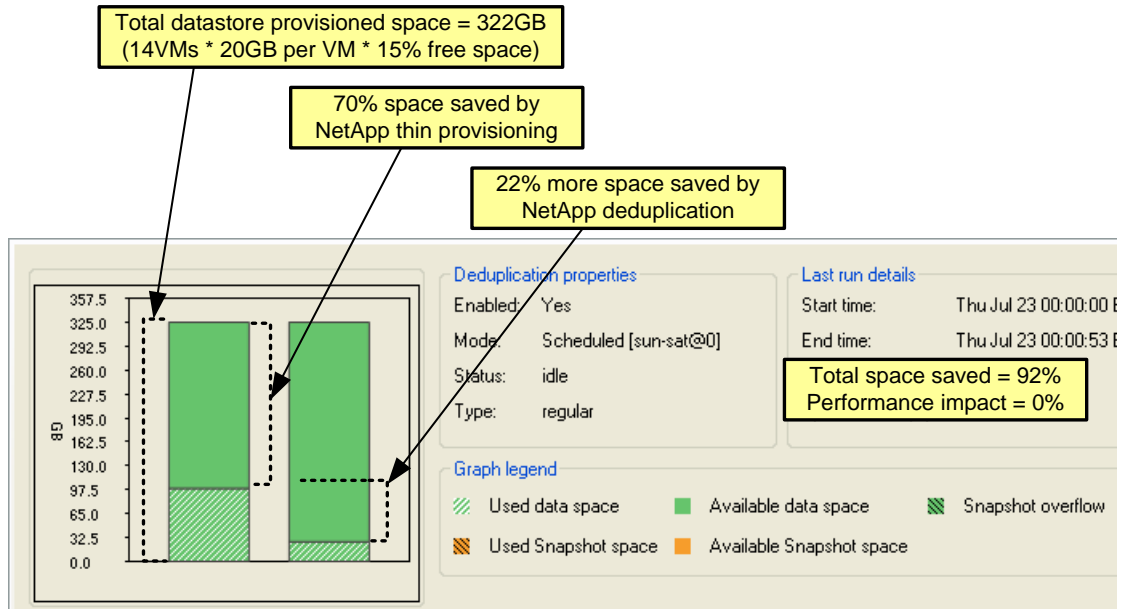


Figure 11) NetApp System Manager showing multiple levels of storage efficiency.

Storage savings for the application-specific data drives, such as SharePoint and SQL database, vary from application to application and environment to environment. For savings specific to each application, see NetApp [TR 3505](#).

4.2 PERFORMANCE VALIDATION

The storage configuration described in this guide was validated by configuring the environment described earlier and then performing performance tests using the application specific tools described in this section. The tests were performed individually for SQL Server, SharePoint, and Exchange and also by running all these applications at the same time. The test results discussed in this section validate that the architecture is capable of handling the mixed workload described earlier.

4.2.1 MICROSOFT EXCHANGE 2007

The Microsoft Exchange [Load Generation Tool](#) was used to simulate the 3000 heavy user environment with 250MB per mailbox. Several 8-hour duration load tests were performed, both with and without NetApp deduplication enabled on the VM C: drives hosting the operating system and Exchange binaries.

LOADGEN USER TASKS LATENCY

Of all the LoadGen user tasks, Send Mail is a very important indicator of the overall performance. For all the 8 hour test runs, the latency numbers were well within the Microsoft recommendation of 1000 ms for acceptable user performance. Also, see the following VMware white paper, which compares the performance of 16,000 heavy user Exchange environment across all the storage protocols (FC, iSCSI, NFS) on NetApp storage.

[VMware vSphere™ 4: Exchange Server® on NFS, iSCSI, and Fibre Channel](#)

VM DISK I/O LATENCY

For all the test cycles, the read and write latencies were well within the Microsoft recommendation listed here: <http://technet.microsoft.com/en-us/library/aa995945.aspx>

VM CPU AND MEMORY UTILIZATION

Each Exchange mailbox server was configured with 9.5GB RAM [2GB + (1500 users per mailbox server * 5MB)] and two virtual CPUs. For the entire 8 hour test cycle, there were no CPU or Memory bottlenecks on the VMs or the ESX host.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire 8-hour test cycle, the NetApp FAS3140 storage controller had more than enough capability to handle the workload for the 3,000 user Exchange environment that was tested. Also, there were no I/O bottlenecks on the storage array.

The NetApp PAM offers significant benefits in an Exchange environment for reducing the solution cost. For more information, see [“Using Performance Acceleration Module in Exchange 2007 Solutions.”](#)

4.2.2 MICROSOFT SQL SERVER 2008

The Microsoft [SQLIOSim](#) utility was used to simulate the database workload described earlier. Several load tests were performed, both with and without deduplication enabled on the VM C: drives hosting the operating system and SQL server binaries.

VM DISK I/O LATENCY

For all the tests, the read and write latencies for the database files were well within the Microsoft recommendations.

VM CPU AND MEMORY UTILIZATION

Each SQL Server VM was configured with 4GB RAM and four virtual CPUs. For the entire duration of the test cycle, there were no CPU or memory bottlenecks on the VMs.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire duration of the test cycles, the NetApp FAS3140 storage controller had sufficient capability to handle the test workload for the SQL environment. Also, there were no I/O bottlenecks on the storage array.

4.2.3 MICROSOFT SHAREPOINT SERVER 2007

Third-party test tools were used to populate and stress test the SharePoint environment described earlier. Several 2-hour load tests were performed with 25% of the users online at any point in time. Tests were conducted both with and without data deduplication enabled on the VM C: drives hosting the operating system and SQL Server binaries.

VM DISK I/O LATENCY

For all the tests, the read and write latencies for the database files were well within the Microsoft recommendations.

VM CPU AND MEMORY UTILIZATION

The Web servers were configured with 2GB RAM and two virtual CPUs, and the Index and database servers were configured with 4GB RAM and four virtual CPUs. For the entire duration of the test cycles, there were no CPU or memory bottlenecks on any of the VMs.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire duration of the test cycles, the NetApp FAS3140 storage controller had sufficient capability to handle the test workload for the SharePoint environment. Also, there were no I/O bottlenecks on the storage array.

As mentioned earlier, the load tests for different applications were also conducted all at the same time. There were no performance bottlenecks on the storage controllers, network, ESX hosts, or VMs.

4.3 VMWARE VMOTION, HA, AND DRS VALIDATION

During the load tests for different applications and also when all the applications were load tested at the same time, VMs were successfully migrated between different ESX hosts using VMotion without any issues. Also, VMware HA and DRS were tested without any issues, demonstrating a high level of solution availability and resource utilization.

4.4 BACKUP AND RESTORE VALIDATION

4.4.1 MICROSOFT EXCHANGE SERVER 2007

Tables 4 and 5 shows the results of the Exchange 2007 backup and restore testing at different levels of granularity. To validate a successful restore of the Exchange 2007 server, several tests can be performed. First, validate the environment by making sure that the log and database files are in the correct places. Then, if the database can be successfully mounted, the environment should be restored properly. Further validation can be done by sending and receiving e-mails, as well as by checking all other functionality, such as calendar and contacts.

Table 4) Exchange 2007 backup test results.

Backup Level	Local Snapshot Backup	SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual datastore	Successful	Successful

Table 5) Exchange 2007 restore test results.

Restore Level	Restore from Local Snapshot Backup	Restore from SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual datastore	Successful	Successful
Individual or multiple mailbox	Successful	Successful

4.4.2 MICROSOFT SQL SERVER 2008

Tables 6 and 7 shows the results of the SQL Server 2008 backup and restore testing at different levels of granularity.

Table 6) SQL Server 2008 backup test results.

Backup Level	Local Snapshot Backup	SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual datastore	Successful	Successful

Table 7) SQL Server 2008 restore test results.

Restore Level	Restore from Local Snapshot Backup	Restore from SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual or multiple VM drives	Successful	Successful
Individual database	Successful	Successful

4.4.3 MICROSOFT SHAREPOINT 2007

Tables 8 and 9 shows the results of the SharePoint 2007 backup and restore testing.

Table 8) SharePoint backup test results.

Backup Level	Local Snapshot backup	SnapMirror Remote Replication
Entire SharePoint environment	Successful	Successful

Backup Level	Local Snapshot backup	SnapMirror Remote Replication
Individual VMs	Successful	Successful

Table 9) SharePoint restore test results.

Restore Level	Restore from Local Snapshot Backup	Restore from SnapMirror Remote Replication
Entire SharePoint environment	Successful	Successful

5 ISCSI REFERENCE ARCHITECTURE DESIGN

5.1 HIGH LEVEL SOLUTION ARCHITECTURE

The iSCSI solution validation involved two VMware ESX hosts running the mixed Microsoft applications workload (described in section 2) with 30 virtual machines hosted on NetApp shared storage. The virtual machine operating system, installed applications, databases, and logs are hosted on NetApp iSCSI based datastores. The solution high availability is achieved by using VMware HA, NetApp active-active controllers and CISCO Nexus 5020 switches. The backup and recovery solution component includes application-consistent point-in-time NetApp snapshot copies with NetApp SMVI, SME, SMSQL, SMOSS, and NetApp SnapMirror replication.

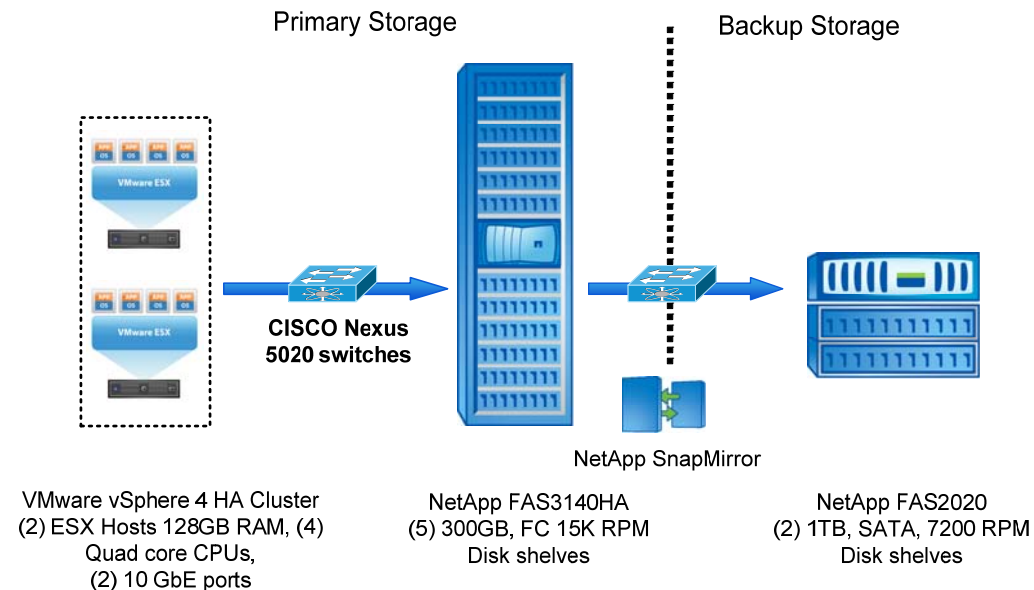


Figure 12) High level solution architecture

5.2 SOLUTION HARDWARE AND SOFTWARE REQUIREMENTS

5.2.1 HARDWARE RESOURCES

Table 10 shows the equipment that was used in this configuration.

Table 10) Hardware configuration.

Solution Component	Minimum Revision
Primary Storage	
(1) NetApp FAS3140HA Cluster	Data ONTAP 7.3.1.1

Solution Component	Minimum Revision
(5) disk shelves	70 disks (14 per shelf); each disk 300GB / 15K/ FC Note: For validating the iSCSI solution, 5 shelves were available as compared to 6 shelves available for the NFS solution. Therefore, there are some differences in how aggregates were used for hosting the VM OS, vswap and page files. This should not be considered a limitation of any protocol.
(2) Cisco Nexus 5020 switches	
(1) Dual port 10Gb Ethernet NIC per FAS3140 controller	
Backup Storage	
(1) NetApp FAS2020HA cluster	Data ONTAP 7.3.1.1
(2) Disk shelves	28 disks (14 per shelf); each disk 1TB/7200RPM/SATA
(2) ESX hosts	
128GB RAM	
(4) Quad Core Xeon Processors	
(1) Dual port 10Gb Ethernet NIC	

5.2.2 SOFTWARE RESOURCES

Table 11 shows the software components that were used in the configuration.

Table 11) Software components.

Solution Component	Minimum Revision
Primary Storage	
Data ONTAP	7.3.1.1
NFS, ASIS, FlexClone, SnapMirror, SnapRestore, and NearStore Licenses	N/A
Backup Storage	
Data ONTAP	7.3.1.1
NFS, ASIS, SnapMirror and NearStore Licenses	N/A
NetApp Management Software	
NetApp Virtual Service Console (VSC)	1.0
NetApp System Manager	1.0
NetApp SnapManager for Virtual Infrastructure (SMVI)	1.2
NetApp SnapManager for Exchange	5.0
NetApp Single Mailbox Recovery for Exchange	5.0

Solution Component	Minimum Revision
NetApp SnapManager for SQL Server	5.0
NetApp SnapManager for SharePoint	2.0.2
NetApp SnapDrive [®]	6.1
NetApp DataFabric Manager (DFM)	3.8
VMware vSphere Infrastructure	
ESX Hosts	VMware ESX , 4.0.0 (build 164009)
vCenter Server	4.0.0
vCenter Database	SQL Server 2005
Applications Virtual Machine Operating System	
Windows Server 2008	x64, Enterprise Edition, Service Pack2
Microsoft Applications	
Microsoft Exchange	Exchange 2007 Enterprise Edition SP1
Microsoft SharePoint	SharePoint Server 2007 Enterprise Edition SP2
SQL Server	SQL Server 2008 Enterprise Edition

5.3 ISCSI SOLUTION ARCHITECTURE DETAILS

5.3.1 VIRTUAL MACHINE LAYOUT

The solution contained a total of 30 virtual machines. The purpose of this configuration was to simulate a real-world customer environment, including the supporting utility and test and dev servers, in addition to the primary Microsoft application servers.

MICROSOFT APPLICATIONS VMS

- **Microsoft Exchange Server 2007.** Total 6 VMs (2 Mailbox servers, 2 HUB servers, 2 CAS servers)
- **Microsoft Office SharePoint Server 2007 .** Total 4 VMs (2 Web front end, 1 search/index, 1 SQL Server 2008)
- **Microsoft SQL Server 2008.** 2 VMs
- **Microsoft IIS.** 4 VMs

TEST AND DEV VMS

- **Windows Server 2008.** 4 VMs

UTILITY VMS

(1) Microsoft WSUS, (1) Microsoft SCOM, (1) Microsoft Exchange LoadGen Tool, (4) SharePoint Test Workstations, (1) NetApp DataFabric Manager, (1) VMware vCenter 4, (1) VMware vCenter 4 database

All the NetApp recommended iSCSI settings highlighted in the NetApp vSphere Best Practices TR 3749 were set directly from the vCenter GUI using the NetApp VSC vCenter plugin.

5.3.2 NETWORK ARCHITECTURE

In this solution, the network comprised of two Cisco Nexus 5020 switches. Since the Cisco Nexus switches used in this configuration support virtual port channeling (vPC), logical separation of the storage network from the rest of the network is achieved while at the same time providing a high level of redundancy, fault

tolerance, and security. With the vPC feature, scalable Layer 2 topologies can be deployed, reducing the dependence on Spanning Tree Protocol for redundancy and loop avoidance. Also, high cross-sectional bandwidth is attained by the features ability to use all available physical links that interconnect the devices.

On the Nexus network ensure the following configurations:

- Be sure to setup a management VLAN for the service console, a public VLAN for the virtual machine network, and a private, non-routable VLAN for VMotion™
- Be sure to use a 10Gb connection between the two Nexus 5020 switches
- Be sure to enable a vPC between the two Nexus 5020 switches. In order to use this feature please be sure to have the Cisco NX-OS Software Release 4.1(3)N1 for Cisco Nexus 5000 series switches installed on your Cisco Nexus 5020

While the Cisco Nexus 5020 switches are 10Gb, they do support 1Gb modules. Therefore, other Cisco switches can be used in conjunction with the Nexus 5020's in order to further scale out the virtualization and storage network.

ESX HOST NETWORK ARCHITECTURE

Figure 13 shows the virtual network layout for each ESX host. Each ESX host has two 10Gb Ethernet ports configured into different port groups as shown below. Please note that there are two VMkernel storage ports as vSphere 4 supports multiple TCP sessions with iSCSI datastores. By enabling multiple TCP sessions with the ESX Round Robin PSP (path selection plugin) allows for iSCSI datastores to send I/O over every available path to the iSCSI target (NetApp storage array). Also note that storage VMkernel port group VM_ISCSI is created for VMs that will mount disk drives directly inside the guest VM using MS iSCSI initiator. E.g. SQL Server, and MS Exchange mailbox server etc.

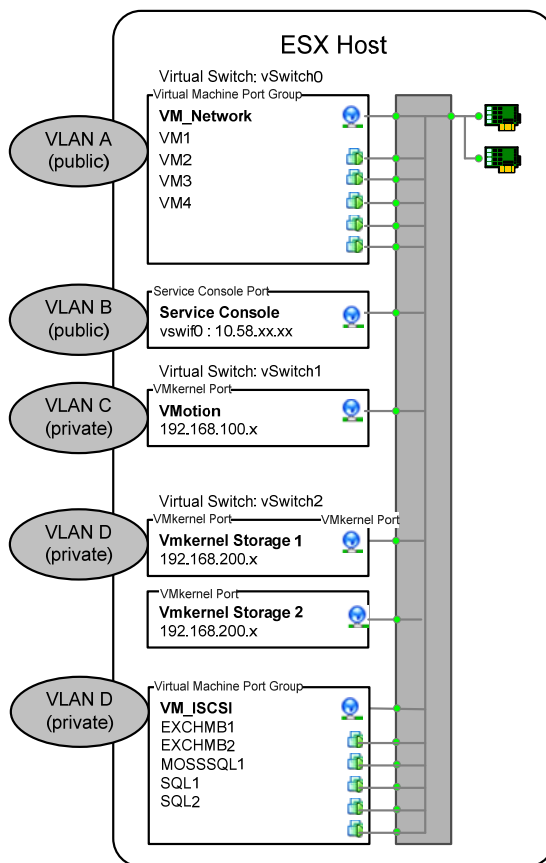


Figure 13) ESX host network architecture

STORAGE NETWORK LAYOUT

Figure 14 shows the storage network layout for the ESX host connectivity with the NetApp storage controller over Cisco Nexus 5020 switches. Make sure to configure a non-routable VLAN for the iSCSI storage traffic to pass to and from the NetApp storage controllers to the vSphere hosts. With this setup the iSCSI traffic is kept completely contained, and security is more tightly controlled.

Also, it is important to have at least two physical Ethernet switches for proper network redundancy in your VMware environment.

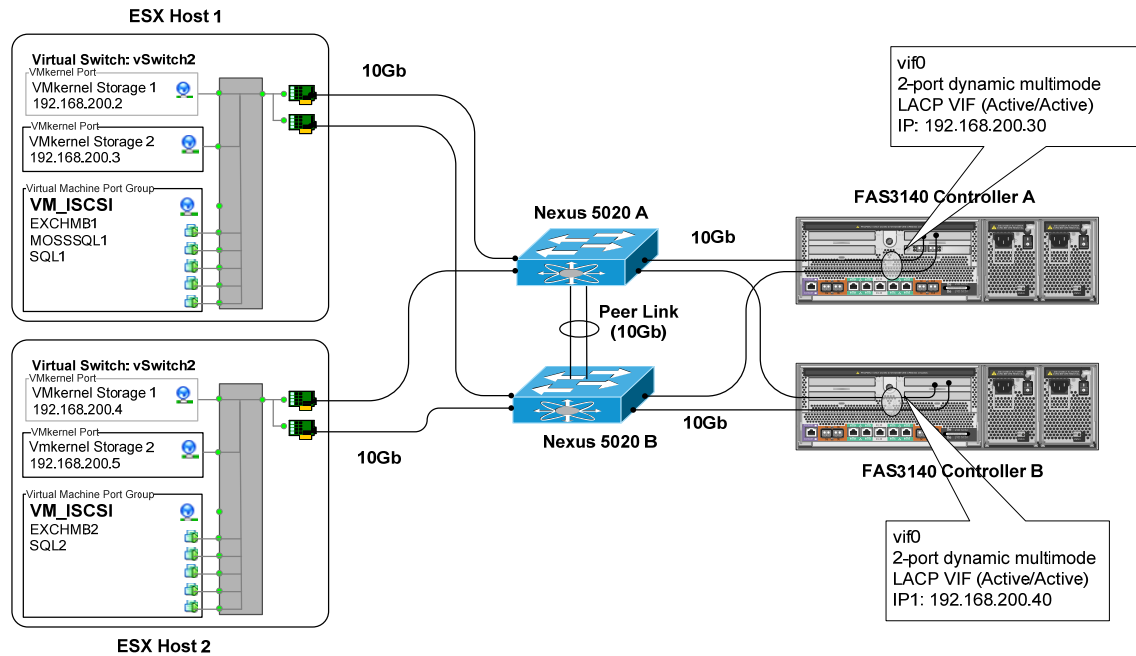


Figure 14) Storage network architecture.

5.3.3 STORAGE ARCHITECTURE

NETAPP STORAGE AGGREGATE LAYOUT

Figure 15 shows the NetApp storage aggregate layout for hosting all the data components for every VM.

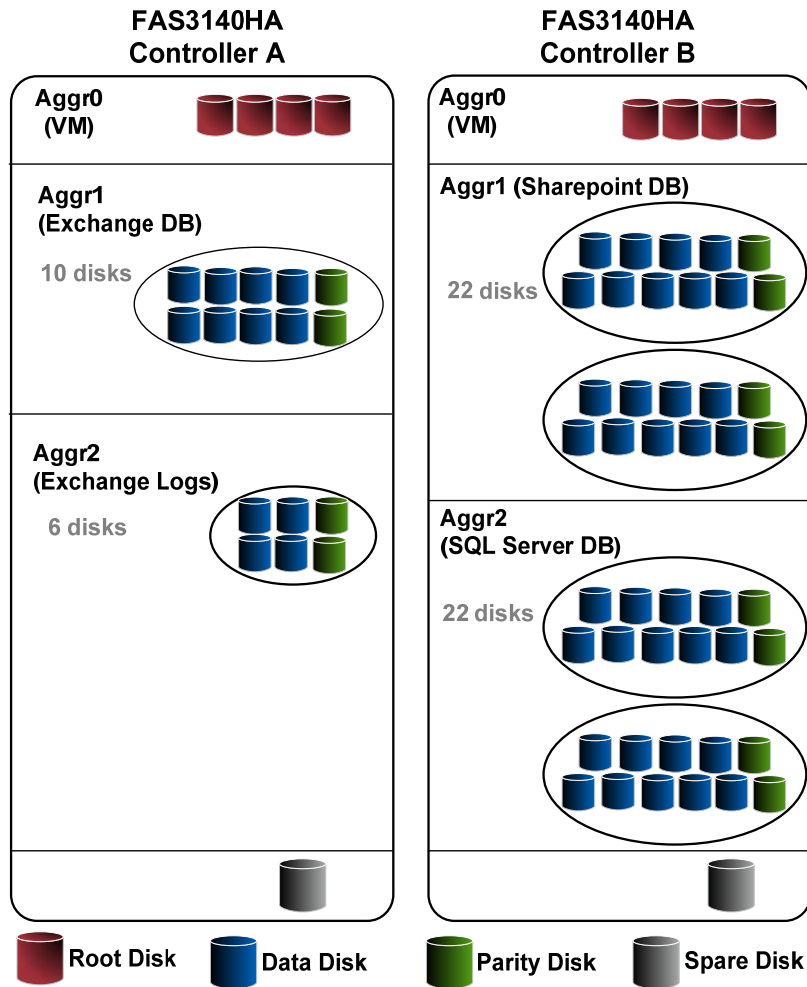


Figure 15) NetApp storage aggregate layout

The aggregate sizing is based on the disk requirements for all the applications to successfully meet their storage capacity and performance requirements.

Note: In this architecture, all the aggregates that host volumes for Exchange VMs are on one storage controller and the aggregates that host volumes for SharePoint and SQL Server are on the second controller. This consideration was made from the perspective of [VMware vCenter Site Recovery Manager](#), which we plan to add in a future release of this guide. VMware vCenter Site Recovery Manager requires all datastores hosting data for a VM to be on the same storage controller.

Also, note that separate aggregates have been created for each application workload. Although this is not required and multiple application workloads can be served off the same aggregate, the reason is to show storage design for each application in a modular way. This will be helpful if you plan to virtualize one of these applications.

NETAPP STORAGE VOLUME LAYOUT

Figure 16 shows the NetApp storage volume layout for hosting the different data components for every VM.

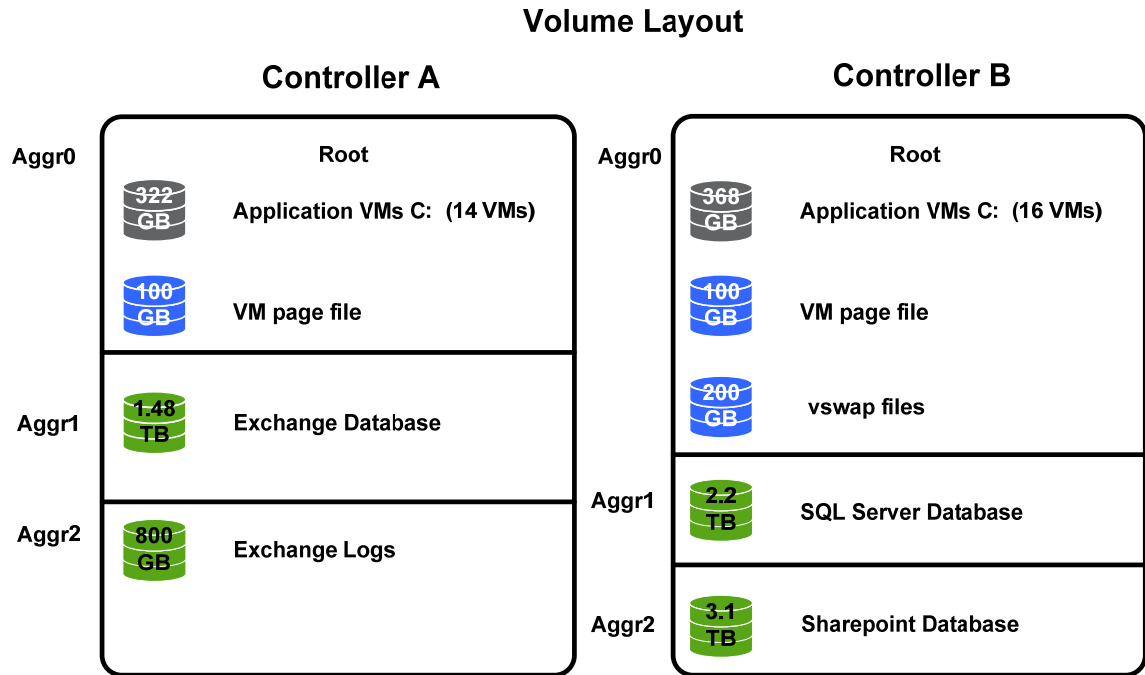


Figure 16) NetApp volume layout

Each virtual machine had a 20GB C: drive with the vmdk hosted on the VM datastore, and a 5 to 10GB T: drive hosted on the page file datastore. In the iSCSI solution, the application servers (Exchange, SQL Server, and SharePoint) database and log drives are hosted on iSCSI LUNs directly connected inside the guest VM using Microsoft iSCSI initiator. This provides the flexibility of leveraging the NetApp and Microsoft applications integrated SnapManager products to achieve granular, automated backup and recovery.

5.3.4 STORAGE SIZING

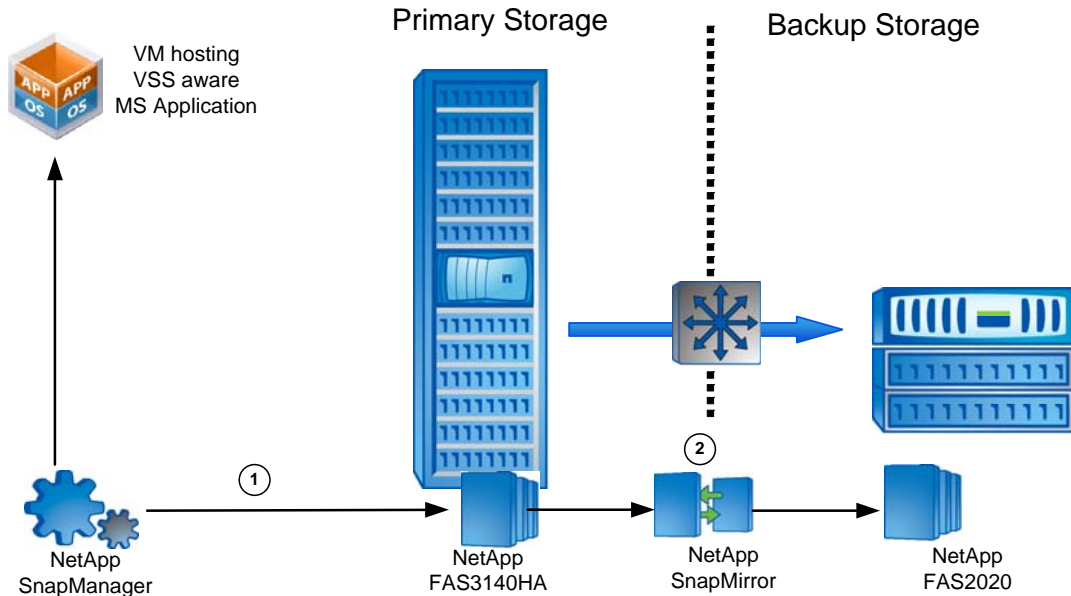
This section contains details about storage sizing. These numbers vary from environment to environment, and you should consult your NetApp systems engineer about the exact sizing for your environment.

- **Microsoft SQL Server 2008.** The SQL Server workload was divided into 10 separate databases using both OLTP and DSS workloads. The databases required 2.2TB of disk space, and the transaction logs required 4.5GB of disk space.
- **Microsoft Office SharePoint Server 2007.** SharePoint used 3TB of disk space for the database and 1GB of disk space for the transaction logs.
- **Microsoft Exchange 2007:** Exchange used 1TB of disk space for the databases and 240GB of disk space for the transaction logs.
- **Datastores hosting VM C: drives (OS and application binaries):** (Size of the VM C: drive + 15% free space for non-vmdk files, VMware snapshots etc.) * Number of VMs in the datastore
 - Application VM datastore on Controller A, hosting 14 VMs = 322GB
 - Application VM datastore on Controller B, hosting 16 VMs = 368GB

These storage requirements are before considering NetApp deduplication and thin provisioning. The NetApp deduplication and storage efficiency savings are highlighted later in the document. The deduplication savings will vary from environment to environment and application to application. For more information, see NetApp [TR3505](#).

5.4 BACKUP AND RESTORE ARCHITECTURE

NetApp SnapManager products were used to provide an application-consistent backup of the database and transaction logs. Remote replication of the backups is achieved using NetApp SnapMirror. Figure 17 shows the backup process workflow.



SnapManager Backup Workflow

1. SnapManager invokes a VSS snapshot of the application data LUNs on the primary storage
2. SnapManager invokes a SnapMirror update to the remote storage controller

Figure 17) High level backup architecture

5.4.1 BACKUP ARCHITECTURE DETAILS

One backup policy is used by SMVI for backing up the datastores hosting the drives with OS and application binaries for the VMs. For Exchange, SQL Server and SharePoint VMs having LUNs directly connected inside guest using MS iSCSI initiator, the VSS driver within the guest OS was disabled. Please refer to the following VMware KB for the justification.

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1009073

For obtaining application consistent backups for these VMs, NetApp SnapManager for Exchange, SQL, and SharePoint were leveraged to perform scheduled backups of the transaction logs and databases.

5.4.2 RESTORE ARCHITECTURE DETAILS

RESTORES FROM LOCALSNAPSHOT BACKUPS

NetApp SnapManager products allow individual application-level recovery. Full VM-level recovery (OS and application) is achieved by using both the NetApp SMVI and application specific SnapManager restore functionality.

RESTORES FROM REMOTE SNAPMIRROR BACKUPS

For a full environment-level restores from the SnapMirror backups on the remote site, follow this process:

1. Quiesce and break the SnapMirror relationships
2. Set up a SnapMirror relationship back from the destination storage system on the remote site to the source storage system at the primary site
3. Quiesce and break the new SnapMirror relationship again

4. Mount the datastores on the ESX host
5. Register the VMs from the restored datastore

Note: Make sure that all the VM hard drives point to the correct vmdk files on the restored datastores.

6 ISCSI REFERENCE ARCHITECTURE SOLUTION VALIDATION

6.1 STORAGE EFFICIENCY

NetApp thin provisioning capabilities were enabled on all the datastores on the primary storage and the deduplication was enabled on the datastores hosting the VM OS and application binaries. The deduplication schedule was set to run once every night. Figure 18 shows screenshot of the NetApp [System Manager](#) that was used to manage the storage, showcasing 92% savings for a datastore hosting OS and Application binaries for 14 virtual machines. Similar storage savings were observed for other datastores on Controller B hosting the OS and application binaries for the SQL Server, and Sharepoint. The saved space (92% in this case) can be used to host additional VMs, additional snapshots, etc. as you scale up the environment.

Please note that as you scale out to virtualize your entire datacenter with 100s to 1000s of VMs, the storage efficiency will be even higher and should be considered in sizing the environment and helping you save cost. Also note that NetApp intelligent caching capabilities (built natively in Data ONTAP and [Performance Accelerator Module](#) cards) strongly complement the NetApp storage efficiency capabilities.

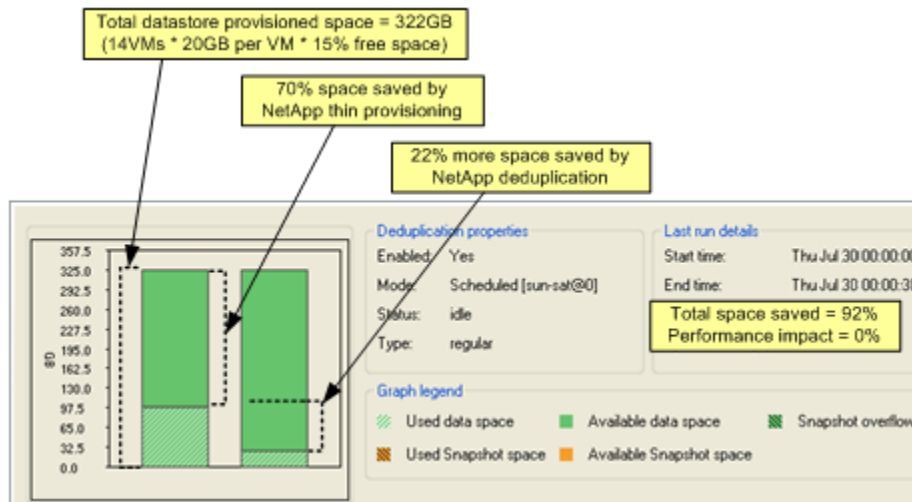


Figure 18) NetApp System Manager showing multiple levels of storage efficiency

Storage savings for the application specific data drives (e.g. Sharepoint, SQL database) will vary from application to application and environment to environment. For savings specific to each application, please refer to the NetApp [TR 3505](#) for further details

6.2 PERFORMANCE VALIDATION

The iSCSI storage configuration described in this guide was validated by configuring the environment described earlier and then performing performance tests using the application specific tools described in this section. The tests were performed individually for SQL Server, SharePoint, and Exchange and also by running all these applications at the same time. The test results validate that the architecture is capable of handling the mixed workload.

6.2.1 MICROSOFT EXCHANGE 2007

The Microsoft Exchange [Load Generation Tool](#) was used to simulate the 3,000 heavy user environment with 250MB per mailbox. Several 8-hour duration load tests were performed, both with and without NetApp deduplication enabled on the VM C: drives hosting the operating system and Exchange binaries.

LOADGEN USER TASKS LATENCIES

Of all the LoadGen user tasks, Send Mail is a very important indicator of the overall performance. For all the 8-hour test runs, the latency numbers were well within the Microsoft recommendation of 1000 ms for acceptable user performance. . Also, see the following VMware white paper, which compares the performance of 16,000 heavy user Exchange environment across all the storage protocols (FC, iSCSI, NFS) on NetApp storage: [VMware vSphere™ 4: Exchange Server® on NFS, iSCSI, and Fibre Channel](#)

VM DISK I/O LATENCY

For all test cycles, the read and write latencies were well within the Microsoft recommendations mentioned here: <http://technet.microsoft.com/en-us/library/aa995945.aspx>

VM CPU AND MEMORY UTILIZATION

Each Exchange mailbox server was configured with 9.5GB RAM [2GB + (1,500 users per mailbox server * 5MB)] and two virtual CPUs. For the entire 8-hour test cycle, there were no CPU or Memory bottlenecks on the VMs or the ESX host.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire 8-hour test cycle, the NetApp FAS3140 storage controller had more than enough capability to handle the workload for the 3,000 user Exchange environment that was tested. Also there were no I/O bottlenecks on the storage array.

The NetApp PAM offers significant benefits in an Exchange environment for reducing the solution cost. For more information, see "[Using Performance Acceleration Module in Exchange 2007 Solutions.](#)"

6.2.2 MICROSOFT SQL SERVER 2008

The Microsoft [SQLIOSim](#) utility was used to simulate the database workload described earlier. Several load tests were performed, both with and without deduplication enabled on the VM C: drives hosting the operating system and SQL Server binaries.

VM DISK I/O LATENCY

For all the tests, the read and write latencies for the database files were well within the Microsoft recommendations.

VM CPU AND MEMORY UTILIZATION

Each SQL Server VM was configured with 4GB RAM and four virtual CPUs. For the entire duration of the test cycle, there were no CPU or memory bottlenecks on the VMs.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire duration of the test cycles, the NetApp FAS3140 storage controller had sufficient capability to handle the test workload for the SQL Server environment. Also, there were no I/O bottlenecks on the storage array.

6.2.3 MICROSOFT SHAREPOINT SERVER 2007

Third-party test tools were used to populate and stress test the SharePoint environment described earlier. Several 2-hour load tests were performed with 25% of the users online at any point in time. Tests were conducted both with and without data deduplication enabled on the VM C: drives hosting the operating system and SQL Server binaries.

VM DISK I/O LATENCY

For all the tests, the read and write latencies for the database files were well within the Microsoft recommendations.

VM CPU AND MEMORY UTILIZATION

The Web servers were configured with 2GB RAM and two virtual CPUs, and the index and database servers were configured with 4GB RAM and four virtual CPUs. For the entire duration of the test cycles, there were no CPU or memory bottlenecks on any of the VMs.

NETAPP STORAGE UTILIZATION SUMMARY

For the entire duration of the test cycles, the NetApp FAS3140 storage controller had sufficient capability to handle the test workload for the SharePoint environment. Also, there were no I/O bottlenecks on the storage array.

As mentioned earlier, the load tests for different applications were also conducted all at the same time. There were no performance bottlenecks on the storage controllers, network, ESX hosts, or VMs.

6.3 VMWARE VMOTION, HA, AND DRS VALIDATION

During the load tests for different applications, and also when all the applications were load tested at the same time, VMs were migrated between different ESX hosts by using VMotion. No issues were observed. Also, VMware HA and DRS were tested without any issues, demonstrating a high level of solution availability and resource utilization.

6.4 BACKUP AND RESTORE VALIDATION

6.4.1 MICROSOFT EXCHANGE SERVER 2007

Tables 12 and 13 shows the results of the backup and restore testing for Exchange 2007 at different levels of granularity.

Table 12) Exchange Server 2007 backup test results.

Backup Level	Local Snapshot Backup	SnapMirror Remote Backups
Entire VM	Successful	Successful
Individual Storage Group	Successful	Successful

Table 13) Exchange Server 2007 restore test results.

Restore Level	Restore from Local Snapshot Backup	Restore from SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual Storage Group	Successful	Successful
Individual Mailbox Recovery (SMBR)	Successful	Successful

6.4.2 MICROSOFT SQL SERVER 2008

Tables 14 and 15 shows the results of the backup and restore testing for SQL Server 2008 at different levels of granularity.

Table 14) SQL Server 2008 backup test results.

Backup Level	Local Snapshot Backup	SnapMirror Remote Backups
Entire VM	Successful	Successful
Individual database	Successful	Successful

Table 15) SQL Server 2008 restore test results.

Restore Level	Restore from Local	Restore from SnapMirror
---------------	--------------------	-------------------------

	Snapshot Backup	Remote Replication
Entire VM	Successful	Successful
Individual database	Successful	Successful
Individual transaction level	Successful	Successful

6.4.3 MICROSOFT OFFICE SHAREPOINT SERVER 2007

Tables 16 and 17 shows the results of the backup and restore testing for Microsoft Office SharePoint Server 2007 at different levels of granularity.

Table 16) SharePoint backup test results.

Backup Level	Local snapshot backup	SnapMirror Remote Backups
Entire SharePoint Environment	Successful	Successful
Individual VMs	Successful	Successful

Table 17) SharePoint restore test results.

Restore Level	Restore from local Snapshot Backup	Restore from SnapMirror Remote Replication
Entire VM	Successful	Successful
Individual database	Successful	Successful
Item level	Successful	Successful

7 SUMMARY

Many customers who are using virtualization today with great success are considering VMware vSphere as the next-generation platform for their Microsoft environment. This report describes how a building-block approach can be used to design an environment with differing disk I/O workloads. This approach leverages a virtualized platform for flexibility and is easily scalable by adding building blocks at any time. Performance testing validated that Microsoft server applications perform extremely well when virtualized on VMware vSphere and NetApp. The integrated NetApp, VMware, and Microsoft backup and recovery NetApp SnapManager tools provided granular, application consistent, space efficient backup and recovery for the entire environment.

VMware VMotion offered significant increase in the solution flexibility for virtual servers and is able to seamlessly move heavily loaded application servers across ESX servers with no loss of service. This proved that VMware VMotion can be a valuable tool for the Microsoft administrator for increased flexibility and avoiding planned downtime.

VMware HA provides a robust solution for protecting every virtual server in the organization from server hardware failure. For customers who have increased availability requirements, VMware HA can be combined with other software clustering solutions such as VMware Fault Tolerance or Windows Failover Clustering.

The architecture as a whole represents an end-to-end solution for deploying Microsoft server applications on a next-generation platform built on VMware vSphere and NetApp IP storage.

8 ACKNOWLEDGEMENTS

The following people have contributed to the Reference Architecture design, validation, and creation of this design guide: Abhinav Joshi, Robert Quimbey, Jack McLeod, Rachel Zhu, John Parker, Vaughn Stewart, Sourav Chakraborty, Amrita Das, Chaffie McKenna, Peter Learmonth, and Abrar Y.

9 FEEDBACK

If you have questions or comments about this document, send email to xdl-vgibutmevmtr@netapp.com

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.