Technical Report

Disaster Recovery of Microsoft Exchange, SQL Server, and SharePoint Server Using VMware vCenter Site Recovery Manager, NetApp SnapManager and SnapMirror, and Cisco Nexus Unified Fabric

Larry Touchette, Sitakanta Chaudhury, NetApp
Abhinav Joshi, Cisco
June 2011 | TR-3822

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 INTRODUCTION

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

Virtualizing Microsoft applications on NetApp unified storage provides enhanced data protection and fully automated disaster recovery capabilities. Typically, a virtualized Microsoft application provides much greater flexibility and complete automation with predefined disaster recovery (DR) processes and easier, streamlined recovery in the event of a disaster.

This solution guide describes the architecture and validation of recovery of Microsoft applications with VMware vCenter™ Site Recovery Manager (SRM) in the NetApp environment. For detailed information on SRM, refer to TR-3671: VMware vCenter Site Recovery Manager in a NetApp Environment.

Some key benefits of the overall solution are:

- **Reduced disaster recovery (DR) costs with SRM.** SRM decreases the risks associated with traditional DR. Repeatable, reliable DR processes are easily developed and maintained. Recovery time objective (RTO) durations are improved through the use of DR workflow automation features. In addition, SRM provides a valuable DR testing capability that allows you to quickly and nondisruptively perform DR tests. Maintaining identical physical hardware at both locations is not required; other processes such as testing and development are enabled to make use of the equipment at the DR site.
- **Reduced infrastructure costs with NetApp storage solutions.** NetApp's disaster recovery solution helps reduce cost so that the savings achieved at the primary site can be used to fund the DR site. With the NetApp solution, you can achieve multiple levels of storage efficiency at the primary site, and the savings are preserved at the DR site. Built-in WAN acceleration strongly complements this because it reduces the amount of WAN bandwidth required for site-to-site replication transfers by more than 80%. NetApp technologies eliminate the need to maintain twice the disk space at the DR site if you want to perform DR testing without interrupting the site-to-site replication. The key NetApp technologies that enable this are thin provisioning, deduplication, NetApp FlexClone®, and NetApp SnapMirror® compression.
- **Application-consistent disaster recovery.** NetApp provides the capability to recover applications in a consistent state after failover to the DR site. The NetApp application-specific SnapManager® products along with SRM offer the ability to maintain a history of multiple, verified, application-consistent recovery points at the DR site. The NetApp SnapManager solution is built using integrated VMware, Microsoft, and NetApp technologies for advanced, application-aware data protection.
- **Simplified disaster recovery processes.** DR workflow automation provided by SRM and the NetApp Storage Recovery Adapter allow the customer to implement testable, repeatable, and, most importantly, reliable disaster recovery processes.
- **Flexible and secure multi-tenancy (SMT).** NetApp unified storage with MultiStore® capability offers end-to-end data security, nondisruptive data mobility, load balancing across storage controllers, and better manageability in a multi-tenant cloud environment. NetApp MultiStore divides a single storage system into multiple secure partitions called vFiler™ units. Individual vFiler units can be assigned to separate "tenants," which can be individual organizations, departments within an organization, or individual applications. The vFiler units can also be shared by multiple organizations depending on the requirements. For more details on the NetApp SMT solution, refer to http://www.netapp.com/us/technology/secure-multi-tenancy.html.

Figure 1) NetApp Multistore with vFiler units.

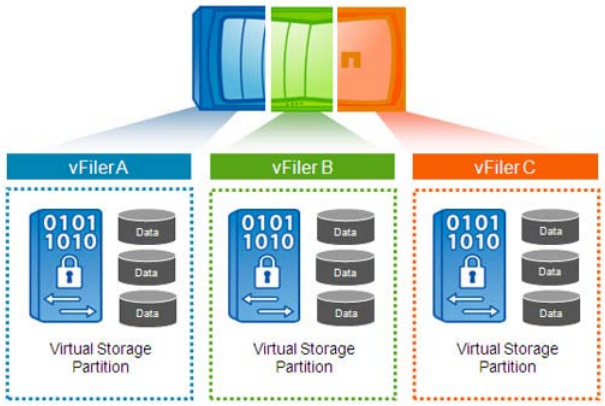

The unique design of MultiStore allows it to support NetApp Data Motion for nondisruptive migration of tenant data between storage systems. With NetApp Data Motion, an entire vFiler unit can be migrated from one storage system to another without disrupting ongoing tenant activity. NetApp Data Motion does for data what VMware VMotion® does for virtual machines (VMs), making it simple to migrate data at a large scale. Combining these services with NetApp Data Motion provides mobility at every layer of your vSphere infrastructure for load balancing, nondisruptive upgrades, or to satisfy other data center needs. For detailed information on MultiStore, refer to http://www.netapp.com/us/communities/tech-ontap/tot-secure-mobile-cloud-storage-1001.html.

2 SOLUTION SUMMARY

This solution guide offers guidance on the design and validation of the SRM and NetApp solution for disaster recovery of the Microsoft applications mixed workload similar to the workload described in TR-3785: Microsoft Exchange Server, SQL Server, and SharePoint Server Mixed Workload on VMware vSphere 4, NetApp Unified Storage (FC, iSCSI, and NFS), and Cisco Nexus Unified Fabric.

In addition to the recovery of the Microsoft Exchange, SQL Server®, and SharePoint® environments, this guide also describes processes for providing the Microsoft Active Directory® and Domain Name System (DNS) services on which the applications are highly dependent.

SRM features such as DR testing, DR failover, and DR workflow automation of Microsoft applications have been successfully validated.

In this solution guide, the virtualized Microsoft applications are protected and recovered by SRM, NetApp SnapMirror, and application-specific SnapManager products.

To verify the validity of this solution design in SMT environments, the processes were validated on both NetApp physical arrays and with Microsoft applications data stored in MultiStore vFiler units. For MultiStore validation, each application was hosted on a different vFiler unit with a corresponding vFiler unit on the SRM recovery site.

Table 1 provides the details of the application VMs tested in this solution.

Table 1) Microsoft application virtual machines.

Microsoft Application	Virtual Machine
Microsoft Exchange Server 2010	Two Exchange mailbox servers
	Two Exchange CAS servers, two hub servers
Microsoft SQL Server 2008 R2	Two SQL Servers
Microsoft SharePoint Server 2010	Two Web/query servers
	One index server, one SQL Server
Microsoft Active Directory and DNS Services	Two AD/DNS servers at protected site and one at recovery site

The key highlights of this solution are:

- Microsoft applications virtualization with VMware vSphere 4
- Microsoft applications DR with SRM and application consistency with NetApp SnapManager and SnapMirror products
- Integrated, storage-efficient DR testing with NetApp FlexClone
- Storage efficiency, without any negative trade-offs, with NetApp primary storage deduplication, FlexClone, and thin provisioning
- Efficient WAN acceleration with deduplication-aware built into NetApp SnapMirror compression
- Solution supportability for SMT environments using NetApp MultiStore vFiler units with Data Motion

NETAPP SNAPMANAGER RECOVERY POINTS

NetApp SnapManager for SQL Server, SnapManager for Exchange, and SnapManager for SharePoint provide flexible DR options by allowing the creation and verification of a history of multiple recovery points that are replicated to the recovery site. The SnapManager applications create full backups, which are verified to be application consistent, and more frequent backups that include only the incremental logs of changes that have occurred between full backups. These incremental backups are referred to as frequent recovery points or FRP backups. Adjusting the time between FRP backups offers the flexibility to set the desired recovery point objective (RPO) for each application separately.

Multiple recovery points provide additional data consistency assurance at the recovery site. If any issues are detected with the recovered application data, individual applications can be reverted to any previously created application-consistent SnapManager recovery point. The SnapManager products also provide the capability to roll forward any uncommitted database logs if the applications are reverted to a previous recovery point. This prevents the loss of any new data that was written at the recovery site after failover.

Figure 2 shows the use of FRP backups to determine the RPO of the solution.

Figure 2) SnapManager recovery points.

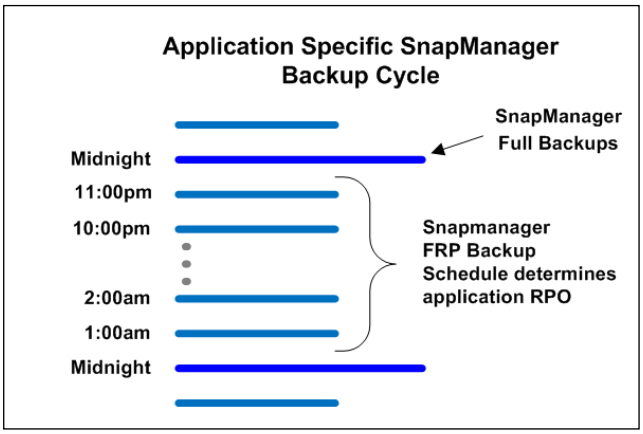

NETAPP STORAGE EFFICIENCY FOR DR TESTING

The NetApp Disaster Recovery Adapter for SRM uses FlexClone to create a writable copy of the data at the recovery site for DR testing. NetApp FlexClone volumes share common data blocks with their parent volumes but behave as independent volumes. There is no need to preprovision containers for clones or to split a clone so that replication can continue. FlexClone requires only a small percentage of additional capacity for FlexClone metadata. Additionally, any new writes that occur during DR testing would require additional capacity.

DR testing may need to be performed multiple times to verify the functionality of custom scripts in the recovery plan (if any), or any time changes that have been made in the primary environment that might affect the success of the recovery plan. With these storage efficiency capabilities, DR testing can be performed any time, even for an extended period of time, while replication is still maintained to the parent FlexVol® volume in the background.

Figure 3 shows the minimum amount of additional storage that was required for DR testing in our solution.

Figure 3) Typical storage requirement for DR testing.

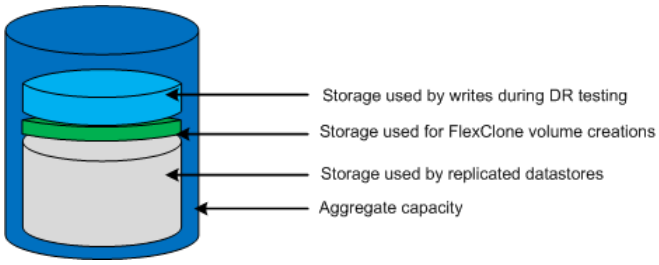

3 SOLUTION DESIGN

3.1 HIGH-LEVEL SOLUTION ARCHITECTURE

The SRM protected site comprises three VMware ESXi 4.1 hosts running the mixed Microsoft applications workload described previously, with VMs hosted on NetApp shared storage. The SRM recovery site is made up of three VMware ESXi 4.1 hosts using NetApp shared storage for SnapMirror destinations. High availability (HA) is achieved at each site by using VMware HA, NetApp active-active controllers, and Cisco® Nexus 5020 switches. The application-consistent data protection solution components include point-in-time NetApp Snapshot™ copies with NetApp SnapManager and NetApp SnapMirror replication as discussed previously and as highlighted in Figure 4.

The validated details of the FC and Ethernet environments are covered in the following subsections.

FC ARCHITECTURE

This solution uses FC protocol for VMFS datastore. SnapDrive 6.3 is used for creating and connecting RDM LUNs. SnapDrive® 6.3 and VSC 2.0.1 or later is used for connecting the virtual machine disks (VMDKs). The FC architecture is valid for environments hosted on NetApp physical arrays.

For each VM, the OS, application binaries, and page files are hosted on VMDK files (hosted on Fibre Channel [FC] VMFS datastore), and the application data is stored as follows:

Case 1: The application server (Exchange Server, SQL Server, and SharePoint) database and log drives are hosted on FC-based raw device mapping (RDM) LUNs that are directly created and connected inside the guest VMs using SnapDrive 6.3.

Case 2: The application server (SQL Server and SharePoint) database and log drives are hosted on FC-based FC VMFS datastores. vCenter Server is used to create virtual machine disks (VMDKs) in FC VMFS datastores that require SnapDrive 6.3 and Virtual Storage Console 2.0.1 or later.

ETHERNET ARCHITECTURE

This solution uses ESXi software iSCSI initiator in ESXi hosts for VMFS datastores. SnapDrive 6.3 is used for creating and connecting RDM LUNs. SnapDrive 6.3 and VSC 2.0.1 or later is used for connecting virtual machine disks (VMDKs).

The Ethernet architecture is valid for environments on NetApp physical arrays, and for Secure Multi-Tenancy environments with VMs and Microsoft application data stored in MultiStore vFiler units.

For each VM, the OS, application binaries, and page files are hosted on VMDK files on Ethernet storage (either iSCSI VMFS datastore or NFS datastore), and the application data is stored as follows:

Case 1: The application server (Exchange, SQL Server, and SharePoint) database and log drives are hosted on iSCSI-based raw device mapping (RDM) LUNs that are directly created and connected inside the guest VMs using SnapDrive 6.3.

Case 2: The application server (SQL Server and SharePoint) database and log drives are hosted on virtual machine disks (VMDKs) in NFS and iSCSI VMFS datastores respectively. vCenter Server is used to create virtual machine disks (VMDKs) in iSCSI VMFS and NFS datastores that require SnapDrive 6.3 and Virtual Storage Console 2.0.1 or later.

Figure 4) High-level solution architecture.

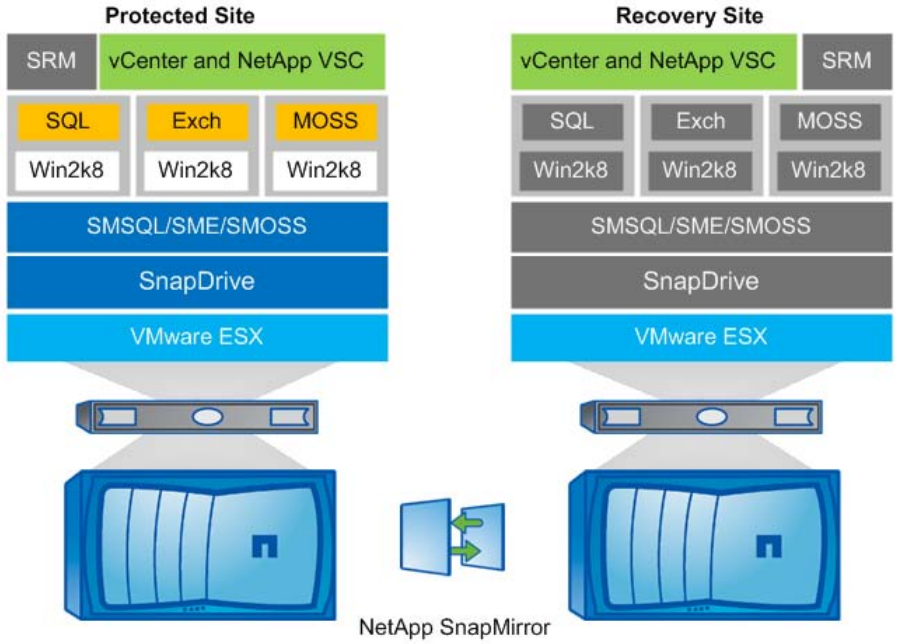

3.2 SOLUTION COMPONENTS

SOLUTION TECHNOLOGY COMPONENTS

- Storage protocols:
  - OS/application binaries: On shared datastores, either VMFS (FC, iSCSI) or NFS
  - Application data: Physical mode RDM LUNs (FC, iSCSI) and Virtual machine disks (VMDKs) on FC, iSCSI and NFS datastores
- VMware: VMware vCenter Site Recovery Manager 4.1, VMware vSphere 4.1
- Cisco: Cisco Nexus Unified Fabric
- NetApp:
  - Remote replication: SnapMirror with built-in compression
  - Application-consistent backups: NetApp Virtual storage console (VSC) , SnapManager for Exchange (SME), SnapManager for SharePoint Server (SMMOSS), SnapManager for SQL Server (SMSQL)
  - Automated provisioning and managing of FC and iSCSI RDM LUNs: NetApp SnapDrive®, leveraging the VMware ESXi S/W iSCSI Initiator

The software components used in the configuration are shown in Table 2.

Table 2) Software components.

Solution Component	Minimum Revision
Storage (Protected and Recovery Sites)	
NetApp Data ONTAP®	8.0.1
FCP or NFS, iSCSI, deduplication, FlexClone, SnapMirror, SnapRestore®, NearStore®, SME, SMSQL, SMMOSS, SnapDrive for Windows®, MultiStore licenses	N/A

Solution Component	Minimum Revision
NetApp Management Software (Protected and Recovery Sites)	
NetApp Virtual Storage console	2.0.1
NetApp SnapManager for Exchange	6.0
NetApp SnapManager for SQL Server	5.1
NetApp SnapManager for SharePoint	6.0
NetApp SnapDrive	6.3
NetApp Disaster Recovery Adapter for SRM	1.4.3
VMware vSphere (Protected and Recovery Sites)	
ESX hosts	VMware ESXi 4.1.0 (build 260247)
vCenter Server	4.1.0
vCenter Database	SQL Server 2005
vCenter Site Recovery Manager	4.1
Applications Virtual Machine Operating System	
Windows Server® 2008 R2	x64, Enterprise Edition
Microsoft Applications	
Microsoft Exchange Server	2010, Enterprise Edition
Microsoft Office SharePoint Server	2010, Enterprise Edition
Microsoft SQL Server	2008 Enterprise Edition, R2

3.3 ACTIVE DIRECTORY/DNS SOLUTION ARCHITECTURE

The configuration consisted of two Microsoft Active Directory/DNS servers at the SRM protected site providing authentication and name resolution services. One Active Directory/DNS server provided those services at the recovery site for that location.

Note: The Active Directory/DNS architecture plays a very important role in the successful failover, DR test, and failback scenarios. If proper procedures are not used, the failover can result in a USN rollback scenario with a corrupt AD database. Refer to http://support.microsoft.com/kb/875495 for more information about the AD issues.

This section describes the procedures that were followed to avoid such situations.

SRM DR TEST SCENARIO

To provide name resolution and user authentication services in the DR Test Network, clone the AD server at the recovery site just prior to running the DR test. Once the cloning is done, before powering on the VM, be sure to connect the cloned AD server to the DR test network. After the AD VM is powered on in the test network, five Flexible Single Master Operations (FSMO) roles in the Active Directory forest must be seized as per the procedure described in the following Microsoft KB: http://support.microsoft.com/kb/255504. The five roles are Schema master, Domain naming master, RID master, PDC emulator, and Infrastructure master.

Note: This procedure can be used in environments with either stretched or different VLANs across sites, as described in section 3.4.

Figure 5 shows seizing of the RID Master FSMO role.

Figure 5) Seizing FSMO roles.

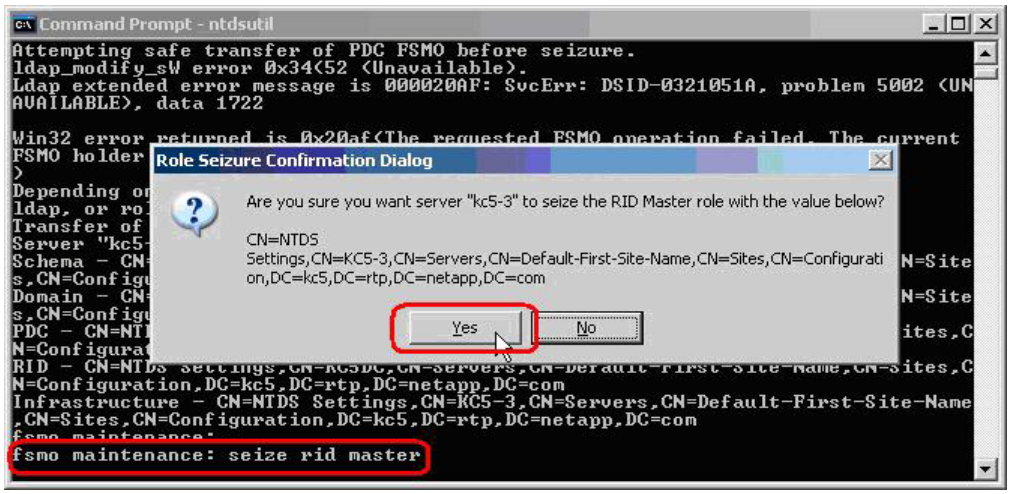

SRM FAILOVER SCENARIO

For a real DR failover scenario, the AD cloning process is not required. In the SRM failover scenario, the existing AD server at the recovery site provides the name resolution and Active Directory services. However, the five FSMO roles must be seized per the procedure described in the Microsoft KB at http://support.microsoft.com/kb/255504.

3.4 MANAGING IP ADDRESS CHANGES BETWEEN SITES

Some environments may be able to use the same network IP addresses at both the protected site and the recovery site. This is referred to as a *stretched VLAN* or a *stretched network setup*. Other environments may have a requirement to use different network IP addresses (in different VLANs) at the recovery site than what is configured at the protected site. In this document, this configuration is referred to as *independent VLANs*. SRM supports both of these scenarios and both have been validated independently in this design.

Microsoft guidelines should be followed so that the Active Directory design for your environment (single site or multisite) is compliant with the business application needs as highlighted in the following Microsoft TechNet document on Microsoft Exchange 2010: http://technet.microsoft.com/en-us/library/aa998636.aspx. This is important for a successful failover of the environment with SRM.

VMware provides a tool called the dr-ip-customizer to assist with configuring SRM to automatically update the VM IP address and the network information (such as primary and secondary DNS server settings) when failover occurs (both the DR test and the real failover). Directions for using the utility are provided in the “SRM Administrator Guide” in the section titled “Customize IP Properties for a Group of Virtual Machines.” The dr-ip-customizer utility is basically a tool that assists with creating a unique customization specification for each VM, and it applies that specification to the recovery plan for each VM.

STRETCHED VLANS

In this scenario, the network IP addressing scheme spans both the protected and the recovery sites. VMs maintain their original IP addresses after failover. In this case, an extra configuration is not needed in SRM provided that the secondary DNS server configured in each VM is accessible at the DR site.

If the primary and secondary DNS servers configured in the VMs are servers that are not available at the DR site after failover, then the VMs must be reconfigured with new primary and secondary DNS server settings. These settings can be applied automatically with the dr-ip-customizer utility.

In this scenario, AD/DNS services can be provided in the private SRM DR test network with either of the two options described in section 3.3.

INDEPENDENT VLANS

In this scenario, the recovery site network requires applying different IP addresses to each VM as it is recovered. This requires changing each VM IP address and updating the appropriate DNS entries for each VM. IP addresses were changed using the dr-ip-customizer utility. The DNS entry for each VM was updated using dynamic DNS. Each VM at the protected site was configured to automatically register its hostname and IP address with the DNS servers configured on the VM. When the VMs are recovered by SRM, their DNS server settings are changed, and the VMs register their new IP addresses with the DNS servers on boot-up.

Note: AD/DNS services in the private SRM DR test network in this design should be provided using only the second option described in section 3.3, by creating a temporary clone of an AD/DNS server at the recovery site and attaching it to the DR test network. The IP address of the domain controller should NOT be changed using the dr-ip-customizer utility, because this process does not account for other steps necessary to correctly change the IP address of a domain controller.

Note: Because the IP addresses of application servers are changed during failover, it is very important to make sure that application servers are not configured to use hard-coded IP addresses to communicate with each other. It is also important to rely on DNS name resolution instead of local hosts or `nbalias` file entries to resolve hostnames into IP addresses, because these files are not updated by the process.

Also make sure that reverse lookup domains are created in DNS, because some applications may depend on reverse lookups. Windows dynamic DNS registration updates both the forward and reverse lookup domains as the VMs register with the DNS server on recovery.

3.5 DR FAILOVER RECOVERY TIME OBJECTIVE

For this type of configuration, the failover times for the entire environment are typically 30 to 60 minutes. The failover times depend on several factors, including the amount of time required for:

- VM, ESXi host, network, and storage-related operations, automated by VMware SRM
- Active Directory operations that are performed manually; seize the five FSMO roles per Microsoft KB http://support.microsoft.com/kb/255504
- Individual application functionality validation after site failover (for example, being able to send and receive e-mails, upload and download documents in SharePoint, common SQL Server operations, and so on)

3.6 NETWORK ARCHITECTURE

Both the protected site and the recovery site are composed of three ESXi hosts. The network architecture in this solution for both the protected site and the recovery site is the same as the solution described in TR-3785.

ETHERNET SOLUTION ARCHITECTURE

The details of the Ethernet architecture based on Cisco Nexus 5020 switches are as follows:

- The two Cisco Nexus 5020 switches had multiple 10GB ports at each site for managing the back-end storage traffic and the IP-based VM Network, VMotion, VMkernel, Management Network, and SnapMirror replication traffic.
- Each ESXi host at both the protected and the recovery sites had two 10Gb Ethernet ports configured with multiple port groups for VM Network, VMotion, VMkernel, and Management Network traffic, as shown in Figure 6.
- For SRM DR testing, one additional virtual machine port group has been added to the networking configuration for the ESXi hosts at the recovery site, as shown in Figure 6. The same is required at the protected site for the failback scenarios.
- Virtual port channeling (vPC) was used on the Cisco Nexus switches to provide 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 bandwidth is attained by the vPC's ability to use all available physical links that interconnect the devices.

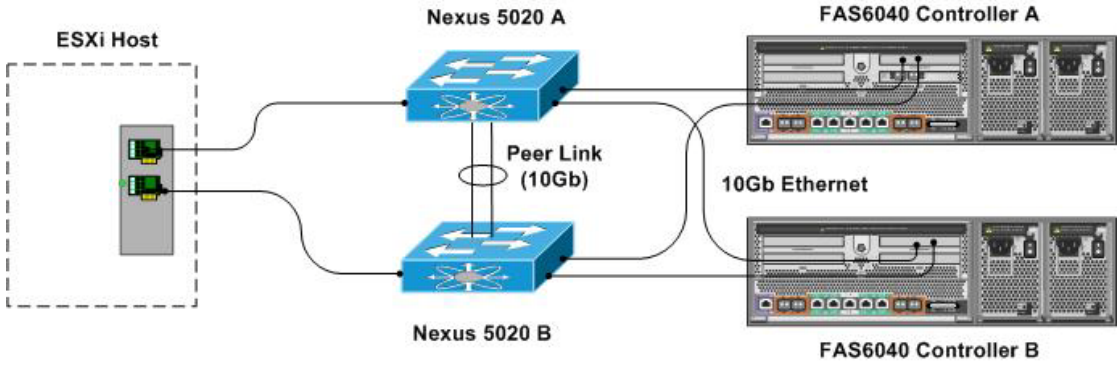

Figure 6) Ethernet solution network architecture.

Make sure that the following configurations are used on the Cisco Nexus network:

- Set up a management VLAN for the management network, a public VLAN for the VM network, and private VLANs for VMotion and VMkernel storage traffic.
- Use a 10Gb connection between the two Cisco Nexus 5020 switches.
- Enable a vPC between the two Cisco Nexus 5020 switches. To use this feature, install the Cisco NX-OS Software Release 4.1(3) N1 for Cisco Nexus 5000 series switches on your Cisco Nexus 5020.

Figure 7) Recovery site ESXi host network diagram for Ethernet architecture (iSCSI, NFS).

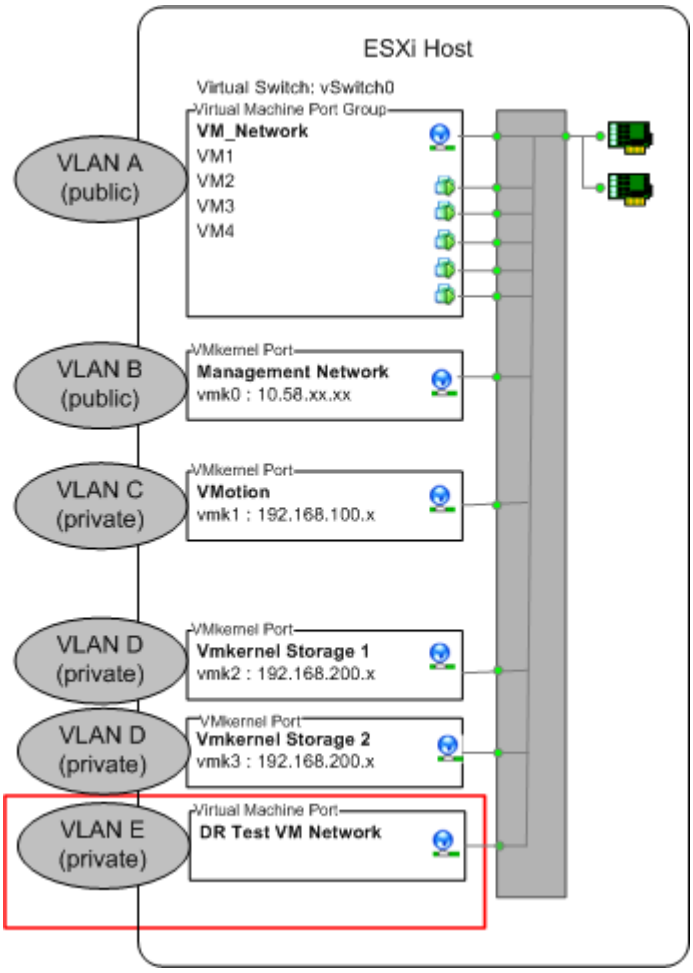

FC SOLUTION ARCHITECTURE

The FC architecture also leveraged the Cisco Nexus 5020 switches with 4GB FC ports configured for storage traffic and 10GB Ethernet ports for the IP-based VM Network, VMotion, Management network, and SnapMirror replication traffic. Each ESXi host had two 4GB FC ports for the storage traffic and two 10GB Ethernet ports for the VM Network, VMotion, and Management network traffic. The ESXi host and network switch configuration for the IP traffic was the same as the configuration used for the previously mentioned Ethernet architecture.

In addition, the Cisco Nexus 5020 switches not only support both the 4Gb FC and 10Gb Ethernet traffic, they also support the 1Gb modules. Therefore, other Cisco switches can be used in conjunction with the Cisco Nexus 5020s to further scale out the virtualization and storage network.

Figure 8) FC Solution network architecture.

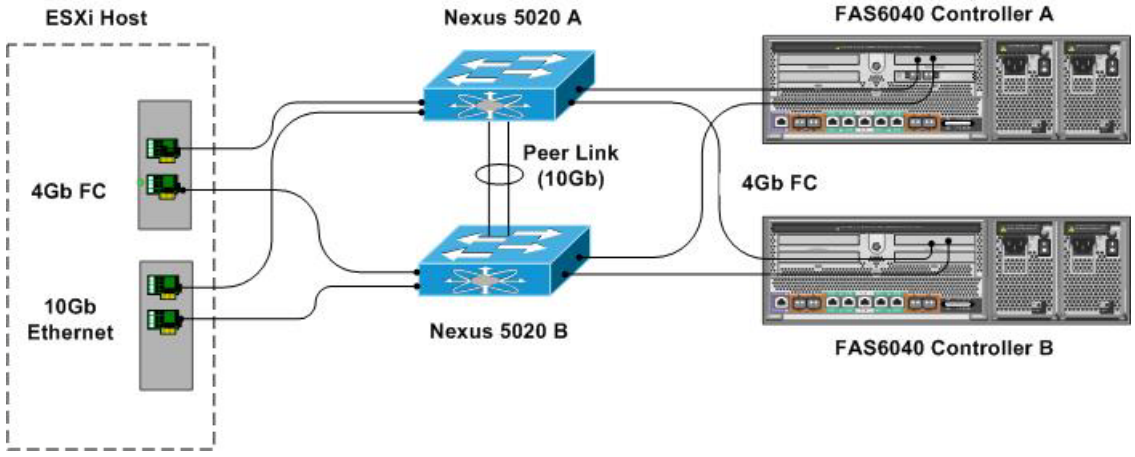

ADDITIONAL SRM NETWORK CONFIGURATION

By default, SRM automatically creates a DR test bubble network on each ESXi host at the DR site and connects VMs to this network instead of to the production networks during DR testing. This is necessary to prevent hostname or IP addressing conflicts on the production networks. However, the automated test bubble networks are created independently on each ESXi host; therefore, VMs recovered on one ESXi host may not be able to communicate with VMs on other ESXi hosts during DR testing. Pre-creating a private DR test network and configuring SRM to use that rather than the automated test network are required for realistic DR testing of multiple application servers across multiple ESXi hosts. The DR test network must be a private physical network or a private VLAN. In this configuration, a private VLAN has been created and will be assigned as the DR test network in the SRM configuration.

3.7 STORAGE ARCHITECTURE

The storage architecture (aggregates, volumes, and datastore layout) is similar to the architecture described in TR-3785.

VIRTUAL MACHINE DATASTORE LAYOUT

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 specifically from the perspective of SRM, which requires all datastores hosting VMDK or RDM data for a VM to be on the same storage controller.

Figure 9 shows the logical storage layout for the different data components of the VMs. The temporary data—VM page file and VM swap file (.vswp)—can be separated into different datastores.

- Separating the VM swap file is easily done with an ESXi cluster setting.
- Separating the VM page file can be done; however, this requires additional steps in SRM (refer to the appendix of TR-3671 for details).

This reduces the daily Snapshot change rate and facilitates faster completion of nightly primary storage deduplication operations and SnapMirror replication. Database and log data is hosted on other volumes.

Figure 9) Virtual machine datastore layout.

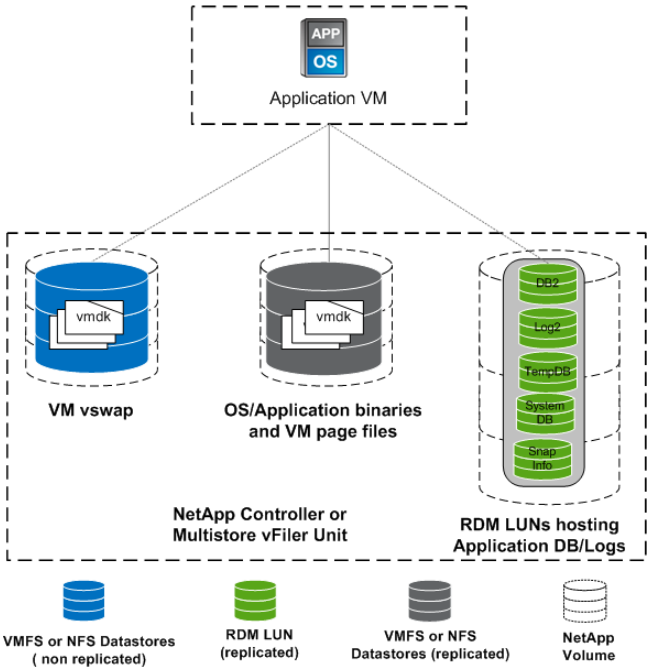

3.8 NETAPP SNAPMIRROR REPLICATION ARCHITECTURE

Data is replicated by using a combination of NetApp backup, recovery, and replication software:

- **NetApp Virtual Storage Console (VSC).** Creates on-disk backups of VMware datastores containing virtual machine system VMDK disks and triggers SnapMirror replication of datastores
- **NetApp SnapManager for Microsoft Exchange.** Creates on-disk backups of Exchange databases and triggers SnapMirror replication of RDM devices containing application data
- **NetApp SnapManager for Microsoft SQL.** Creates on-disk backups of SQL Server databases and triggers SnapMirror replication of RDM and virtual machine disks (VMDKs) devices containing application data
- **NetApp SnapManager for Microsoft SharePoint Server.** Creates on-disk backups of SharePoint databases and triggers SnapMirror replication of RDM and virtual machine disks (VMDKs) devices containing application data

A backup policy in NetApp VSC is used to back up and replicate the VMFS or NFS datastores hosting VMDK files containing the OS and application binaries. In the RDM-based solution (Exchange, SQL Server and SharePoint) and in the VMDKs-based solution (SQL Server and SharePoint), application-consistent backups of data in the Exchange, SQL Server, and SharePoint VMs are created by using NetApp SnapManager for Exchange, SQL Server, and SharePoint. These applications perform scheduled backups of the transaction logs and databases and also initiate SnapMirror updates. The SnapManager products also provide granular recovery points for these Microsoft applications. NetApp highly recommends scheduling the VSC and application-specific SnapManager backups so that they happen at different times.

Figure 10) High-level replication architecture.

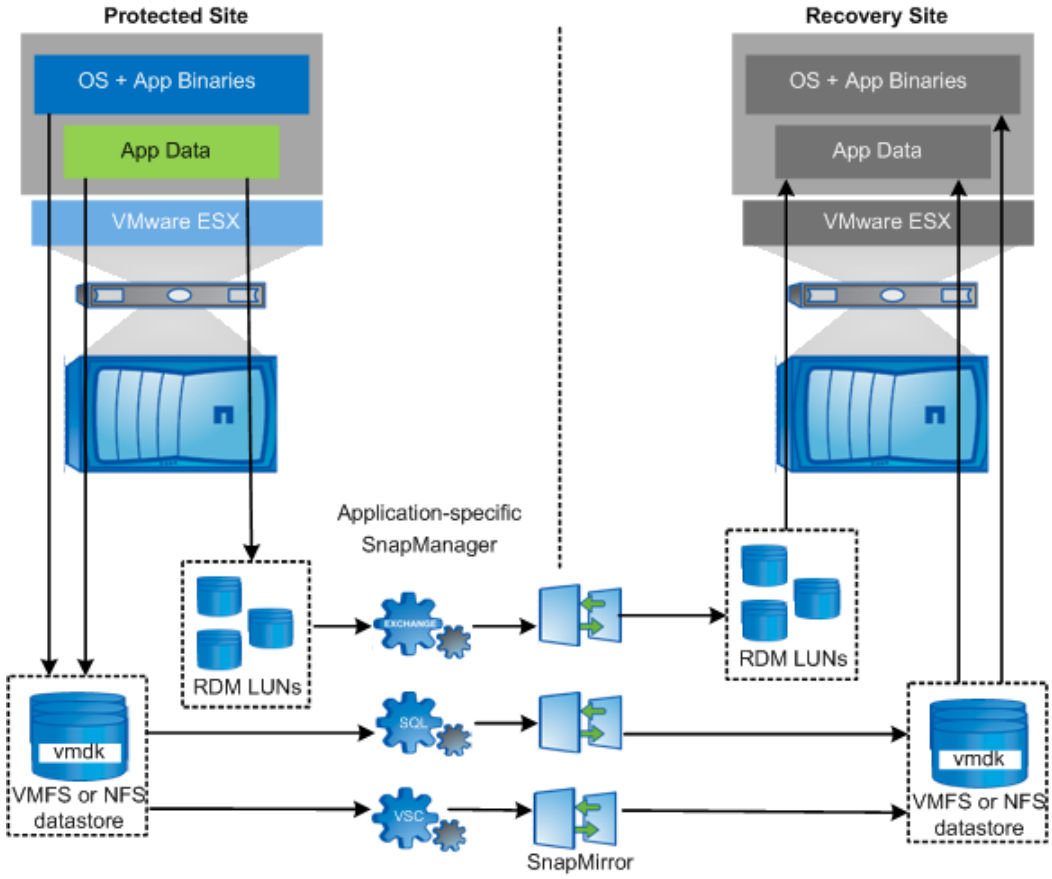

WAN ACCELERATION

High levels of WAN bandwidth savings were achieved by enabling the compression capabilities natively available in NetApp SnapMirror. For a datastore at the protected site hosting Exchange, and SharePoint VMs OS and application binaries, more than 78% storage savings were achieved with NetApp deduplication. A compression ratio of 3:1 was observed during the SnapMirror baseline transfer of this datastore to the recovery site. Therefore, a total savings of 92% on WAN bandwidth was achieved. With daily incremental replication, even higher levels of compression ratios were achieved. For Exchange, SQL Server, and SharePoint application data, compression savings vary based on the data. In our testing for Exchange data, we observed a compression ratio of 1.5:1 and higher.

3.9 SRM CONFIGURATION

SRM PROTECTION GROUPS AND RECOVERY PLANS

As shown in Table 3, protection groups were created in the vCenter server on the protected site based on the group of VMs hosted on the same datastore. One protection group was created for SQL Server VMs and one protection group was created for the Exchange and SharePoint VMs, as described in section 3.3.

Table 3) SRM protection groups and recovery plans.

Protection Group and Recovery Plans	Virtual Machines	VMs in High Priority Step	VMs in Low-Priority Step
Exchange_SharePoint	(2) Exchange Mailbox servers (2) Exchange CAS servers (2) Exchange HUB servers (2) SharePoint Web/query servers (1) SharePoint index server (1) SharePoint database server	(2) Exchange mailbox servers (1) SharePoint index server (1) SharePoint database server	(2) Exchange CAS servers (2) Exchange HUB servers (2) SharePoint Web/query servers
SQL Server	(2) SQL Servers	(2) SQL Servers	N/A

The corresponding recovery plans were created on the vCenter server at the recovery site. One recovery plan was created for the SQL Server VMs, and a second recovery plan was created for the Exchange and SharePoint VMs. For the Exchange_Sharepoint recovery plan, the two Exchange mailbox servers were added in the high-priority VM step so that they power on first and are ready before the Exchange HUB and CAS servers are powered on in the low-priority VM step. The SharePoint database and index servers were added in the high-priority VM step and the Web servers were added in the low-priority VM step so that the database and index servers are ready before the Web servers come up.

Figure 11) Recovery plan VM startup priority.

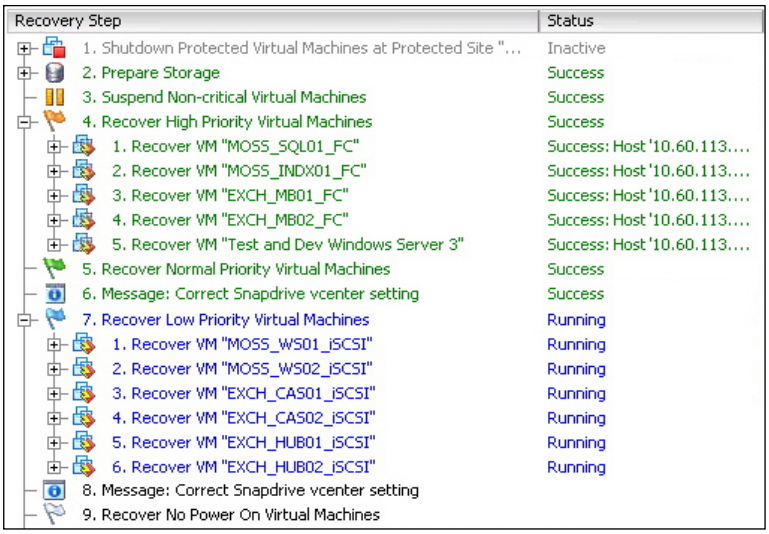

AUTOMATED SNAPDRIVE RECONFIGURATION

NetApp SnapManager for Exchange, SQL Server, and SharePoint applications create on-disk backups of applications. The NetApp SnapDrive software provides virtual machine disks (VMDKs), RDM LUN management and Snapshot capability to SnapManager from within the guest VM operating system.

In a VMware environment, SnapDrive for Windows is configured to communicate with the vCenter server that manages the ESXi host where the VM is running. Through the VMware vCenter server and NetApp VSC, SnapDrive is able to perform management operations on LUNs and VMDKs, such as creating and deleting NetApp Snapshot copies and connecting and disconnecting Snapshot copies of VMDKs and RDM LUNs in the VM. When a VMDK or an RDM LUN is connected to a VM, SnapDrive updates the VM configuration to store the information about that VMDKs or RDM LUN connection. When a failover is executed, the VMDK or RDM LUNs are automatically connected to the VMs upon recovery. Connecting the LUNs or VMDKs to the VM is performed by SRM and not by the SnapDrive software in the guest. After SRM recovery, the LUNs are accessible by the guest OS and applications without intervention by the system administrator. However, to allow management of RDM LUNs or VMDKs by the application-specific SnapManager products while at the DR site, the VMware vCenter server and NetApp VSC settings must be updated in SnapDrive to point to the vCenter server at the recovery site.

Figure 12) SnapDrive communication before and after SRM failover.

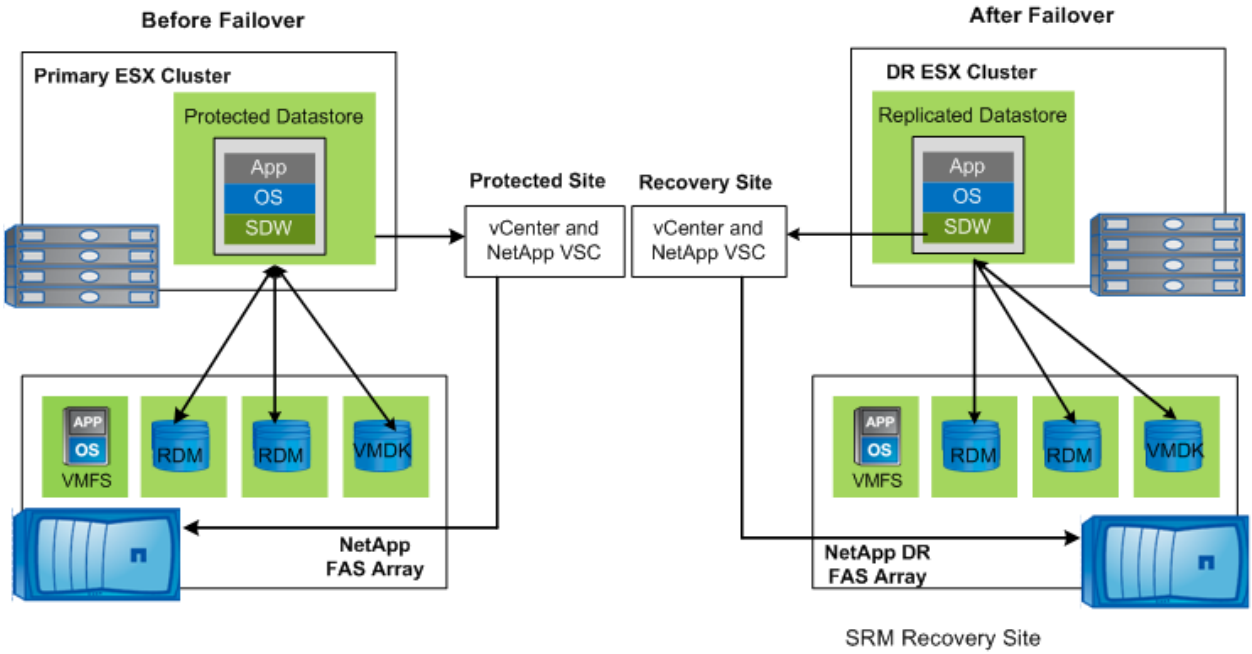

The reconfiguration of the VMware vCenter server settings in SnapDrive can be automated by using command steps in the recovery plans. The command steps are executed on the SRM host. Therefore, when using commands steps, a process capable of sending commands remotely must be used. The SnapDrive command line interface (SDCLI) has this capability. The SnapDrive CLI command `sdcli vsconfig set` can be used to automate the reconfiguration of the SnapDrive vCenter server setting when a failover is performed. For information about the `sdcli vsconfig set` command refer to the "SnapDrive for Windows Administrator Guide."

Requirements for automating the SnapDrive reconfiguration during SRM failover are:

- SRM command steps are executed on the SRM host; therefore, the SnapDrive management software must be installed on the SRM server at both sites.
- In the SRM servers at both sites, the service account that runs the SRM service must be set to an authenticated account that has access to run the `sdcli` command. Otherwise, the `sdcli`

`vsconfig set` command fails. By default, the SRM service runs as a local system, which has no access to the network for executing the `sdcli` command.

3.10 NETAPP VSC RECOVERY CONSIDERATIONS

- After SRM failover, the inventory on the NetApp Virtual Storage Console (VSC) server at the recovery site must be updated to reflect the new datastore identities. For detailed steps, refer to the TR-3737: SMVI Best Practices, which discusses the loss of a primary site.
- This process is required only if the VMs will operate from the recovery site for long periods of time and if backups for OS and application binaries are required. For offsite backups, remote replications must be set up before VSC can continue to replicate the backups to the new remote site or to the original protected site. If the replication must be set up to the original site, only the existing SnapMirror relationships are required to be reversed.
- During SRM DR testing, VSC cannot be used because the vCenter server is inaccessible in the test network. This is true for any backup/DR vendor solution that depends on access to the vCenter server in the DR test mode.
- This recovery process is not required for the application-specific SnapManager products.

4 SOLUTION VALIDATION

4.1 SRM DR TEST PROCESS

Figure 13 shows the high-level SRM DR test process.

Figure 13) DR test summary steps.

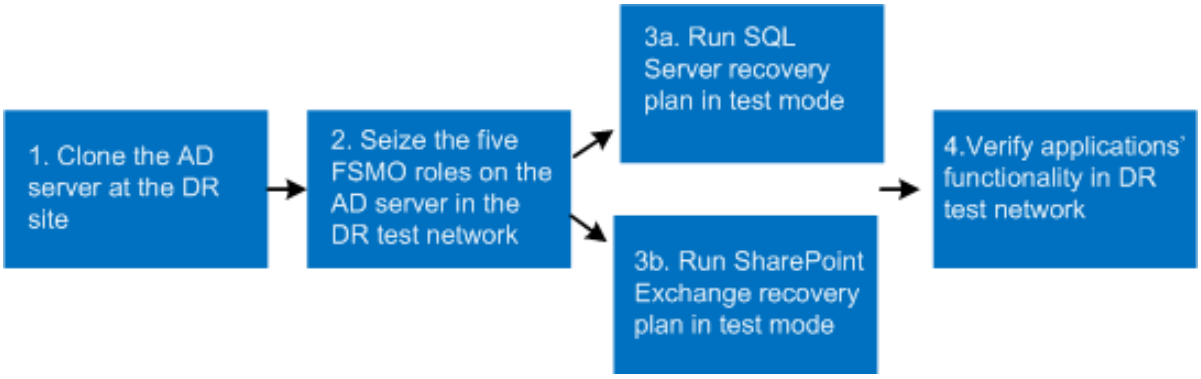

1. **Clone the AD server at the DR site.** To provide AD and DNS services in the test network, clone the AD server at the recovery site just prior to running the DR test. Once the cloning is done, before powering on the VM, be sure to connect the cloned AD server to the DR test network. After the AD VM is powered on in the test network, five FSMO roles in the Active Directory forest must be seized as per the procedure described in the following Microsoft KB: http://support.microsoft.com/kb/255504. The five roles are Schema master, Domain naming master, RID master, PDC emulator, and Infrastructure master.
2. **Seize the FSMO roles on the DR test AD server.** FSMO roles are performed by AD servers in an Active Directory environment. Some of these roles provide services required by the application servers. Before performing the application recovery in the DR test network, an AD server must be providing these services. To provide these, the FSMO roles are seized on the AD server in the recovery site that was cloned and attached to the DR test network. Because this is a private network, this AD server can safely seize these roles to temporarily provide these services in the DR test network without impacting the production AD environment. When the test is complete, the entire test environment is destroyed. The cloned AD server at the recovery site can then be destroyed as well.

Note: Do not connect the AD server holding the FSMO roles in the DR test network to the production network because this will cause conflicts with the production AD servers.

3. **Run application server recovery plans.** Run the recovery plans for each application server. The application servers will be connected to the DR test network and started up.
4. **Verify application functionality.** After the VMs have been powered on in the DR test network, use the vSphere client console connection to access the VMs and verify application functionality.

While in DR test mode, the application VMs have no access to the vCenter server; therefore, the NetApp SnapDrive software cannot be used for SnapManager specific testing. Testing is limited to verifying the application functionality only. This is true for any backup/DR solution that depends on access to the vCenter server while running in DR test mode.

Figure 14) DR test mode.

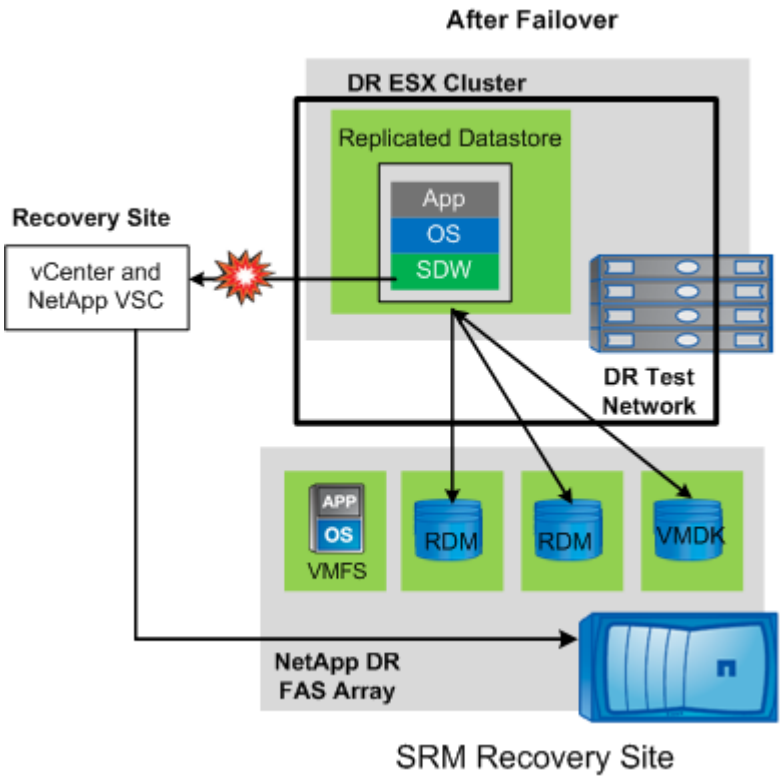

SRM DR TEST RESULTS

Table 4 shows the results of the SRM DR test.

Table 4) DR test results.

Validation Test	Result
Active Directory	
User authentication	✓
Name resolution	✓
Exchange Server	
Send/receive e-mails	✓
SQL Server	

Validation Test	Result
Create/delete tables	✓
Create/delete rows	✓
SharePoint Server	
Upload/download/create documents	✓
Create/delete sites	✓

4.2 SRM FAILOVER PROCESS

Figure 15 shows the high-level SRM DR failover process.

Figure 15) Failover summary steps.

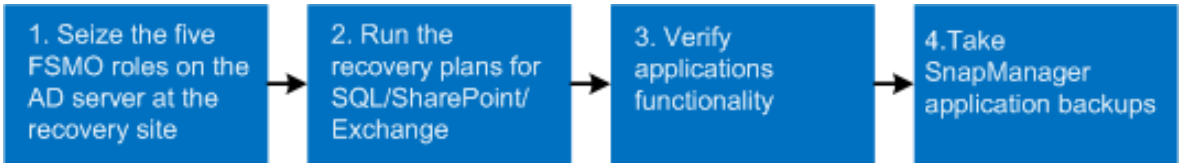

1. **Seize FSMO roles on the AD server at the recovery site.** Before performing the application recovery at the recovery site, an AD server must perform the previously discussed FSMO roles. The recovery of AD servers is not included in the real SRM failover to avoid issues caused by recovering AD servers that have databases that are out of sync with each other, referred to as a USN rollback. Instead, the FSMO roles are seized on an existing AD server at the recovery site.
2. **Run the SRM recovery plans.** After the FSMO roles have been added to the recovery site, run the SRM recovery plans for the SQL Server, SharePoint, and Exchange mailbox servers. In this environment, the Exchange Mailbox servers are recovered first and verified to be available before the Exchange Hub and CAS servers are allowed to start up.
3. **Verify application functionality.** Perform tests to make sure the applications are functional.
4. **Take SnapManager application backup.** SRM and the NetApp Site Recovery Adapter (SRA) recover each VM, automatically connecting each RDM and VMDKs to the VMs, and the applications start up. The applications are recovered to the last NetApp SnapMirror replication point. Once the applications are recovered, a full backup should be created by using the application-specific SnapManager products. These backups should be verified as soon as possible or scheduled to be verified at a later time. If the database verification fails, a SnapManager up-to-the-minute recovery can be performed from the last verified backup. During the SnapManager recovery, all the transaction logs can be rolled forward, including those written at the recovery site after failover. With the SnapDrive vCenter server setting reconfigured, SnapManager automatically continues to create backups on the same schedule as it was configured to do at the protected site. Figure 16 shows the timeline of a recovery scenario after a disaster that occurred at 10:30 a.m.

Figure 16) SnapManager recovery timeline.

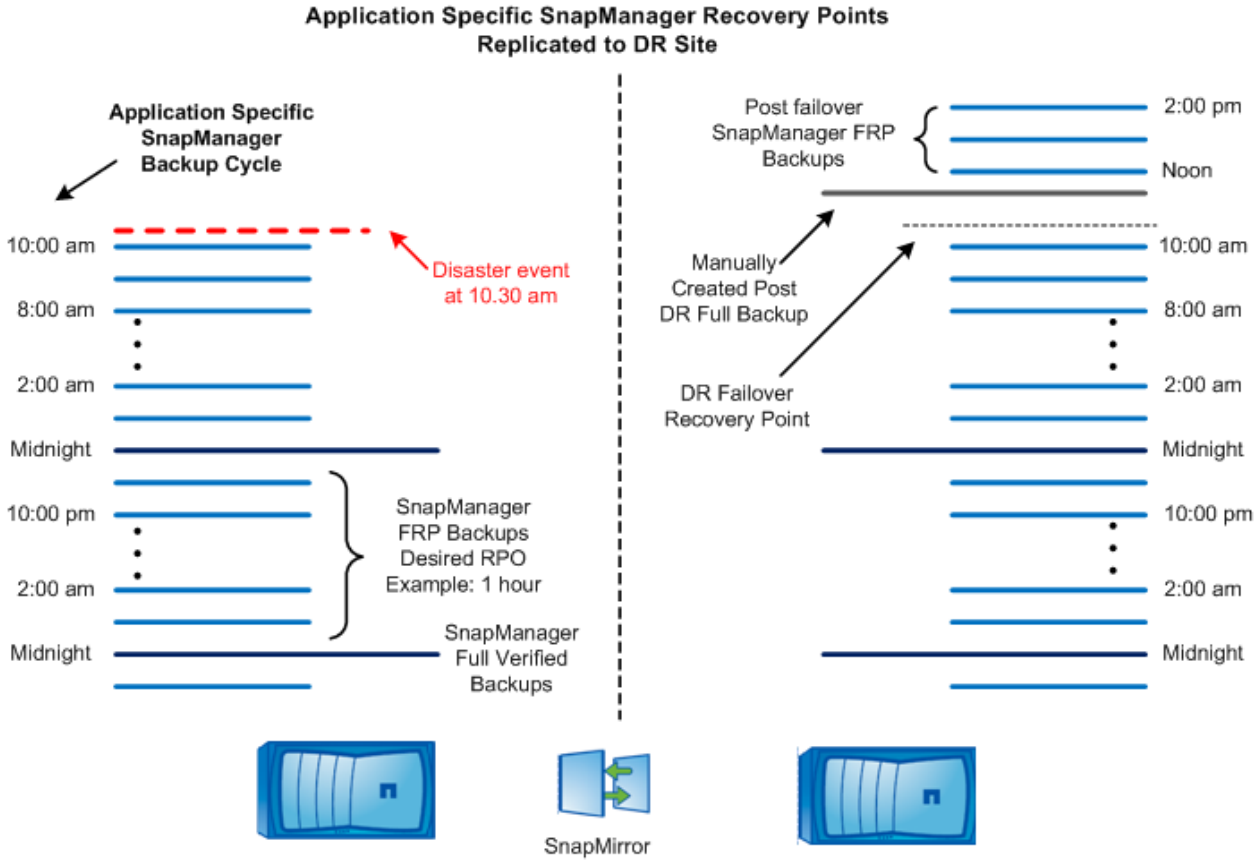

In Figure 16, SnapManager was configured to create full, verified backups of the application on a daily basis at midnight and to create frequent recovery point (FRP) backups every hour, replicating each backup to the recovery site. FRP backups are backups consisting only of transaction logs. The frequency of the FRP backups determines the RPO of the solution. In this example, the RPO for this application server was one hour.

SRM FAILOVER TEST RESULTS

Table 5 shows the results of the SRM failover tests.

Table 5) Failover test results.

Validation Test	Results
Active Directory	
User authentication	✓
Name resolution	✓
Exchange Server	
Send/receive e-mails	✓
SQL Server	
Create/delete tables	✓
Create/delete rows	✓

Validation Test	Results
SharePoint Server	
Upload/download/create documents	✓
Create/delete sites	✓

4.3 SRM FAILBACK PROCESS

Figure 17 shows the high-level failback process.

Figure 17) Failback summary steps.

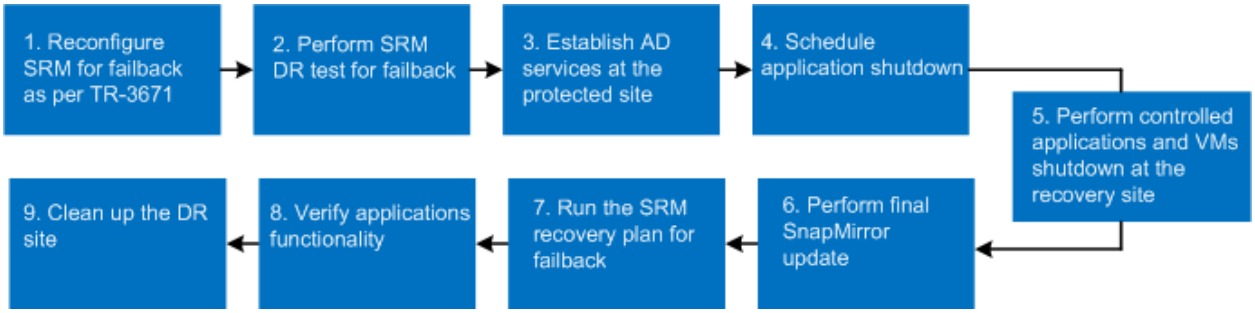

5. **Reconfigure SRM for failback.** Currently, SRM does not provide automatic failback functionality. However, SRM can be reconfigured with protection groups and recovery plans in the reverse direction and used to failback the environment to the original site in a controlled manner. The process for reconfiguring SRM for failback is described in TR-3671. Include the process described in section 3.9 for implementing the SRM command steps to configure the SnapDrive software in the VMs in order to use the vCenter server at the original protected site.
6. **Perform a SRM DR test prior to scheduled failback.** The ability to perform a test prior to failover is as important for failing back to the original site as it was for DR. Using SRM for failback in this way allows a safe and reliable recovery that can be tested prior to a scheduled failback event. Perform a DR test in SRM by using the same method used prior to the disaster.
7. **Establish Active Directory services at the original site.** Prior to failback, Active Directory services should be reestablished at the original protected site. This can be done by recovering the AD servers at that site and forcing them to resynchronize with the newer AD servers at the recovery site or by establishing new AD servers to synchronize. AD servers should not be recovered from replicas created by unsupported processes because this could create a USN rollback scenario. Refer to MS KB888794 for information regarding supported methods of creating backups of AD servers in virtual server environments.

Note: The AD server from which the FSMO roles were seized must not be brought back online at the original protected site. Because that AD server was down at the time the FSMO roles were seized, that server is unaware that it is no longer responsible for providing those services. This server must be retired or recommissioned as a new AD server.

8. **Schedule application shutdown.** Schedule a time for an outage in which the failback can be performed in a controlled manner.
9. **Shut down applications and VMs at the recovery site.** So that all changes made to production data are replicated back to the original site, perform a controlled, orderly shutdown of the application servers and VMs at the DR site prior to the failback.
10. **Perform SnapMirror update.** A SnapMirror update performed while the applications are down makes sure that any new data written just before the application shutdown is replicated to the original protected site.

11. **Run the SRM recovery plans for failback.** When the new recovery plans are run, the VMs are started from the clean shutdown state. The applications are started up, and previously created SnapManager backups of the applications are accessible.
12. **Verify application functionality.** Test the applications as appropriate so that they are available for use.
13. **Clean up the DR site.** SRM does not remove VMs and datastores from the environment at the DR site. This must be done manually. Multiple powered-off VMs may be selected and removed from the vCenter inventory, and then replicated LUNs may be removed from the igroups on the storage arrays. A storage rescan in the vSphere client at the DR site then removes the datastores from the ESXi inventories. The SnapMirror relationships must be reversed and resynced to allow the SnapManager applications to replicate backups to the DR site.

SRM FAILBACK TEST RESULTS

Table 6 shows the results of the failback tests.

Table 6) Failback test results.

Validation Test	Results
Active Directory	
User authentication	✓
Name resolution	✓
Exchange Server	
Send/receive e-mails	✓
SQL Server	
Create/delete tables	✓
Create/delete rows	✓
SharePoint Server	
Upload/download/create documents	✓
Create/delete sites	✓

5 SUMMARY

To summarize, many customers who are using virtualization today with great success are considering VMware vSphere as the next-generation platform for their Microsoft environment. A virtualized Microsoft applications environment provides much greater flexibility and complete automation with predefined DR processes and easier, streamlined recovery in the event of disaster. This solution guide provides guidance on how to successfully architect and validate a fully automated and cost-effective disaster recovery solution for Microsoft applications by using VMware SRM and NetApp software, providing the smoothest possible recovery from disaster.

The solution architecture, setup, and validation concluded that SRM offers the capability to create repeatable, reliable DR processes and helps lower RTO for mission-critical Microsoft applications by providing DR workflow automation capabilities. It also provides DR testing capabilities, allowing customers to quickly and nondisruptively perform DR tests.

NetApp strongly complements the VMware Site Recovery Manager solution for Microsoft applications by providing the most cost-effective and tightly integrated storage solution with VMware vSphere 4, SRM, and virtualized Microsoft applications. The key highlights of the NetApp value-add are:

- Multiple levels of storage efficiency enhance the SRM solution. NetApp's thin provisioning and data deduplication benefits are transferred to the recovery site. In addition, deduplication-aware NetApp

SnapMirror along with built-in network compression reduce the amount of WAN bandwidth required for replication by more than 80%. NetApp FlexClone technology eliminates the need to maintain twice the disk space at the DR site to perform DR testing without interrupting site-to-site replication.

- NetApp SnapManager, along with SRM, provides the capability to achieve application-consistent disaster recovery at the DR site. The NetApp SnapManager solution is built by using integrated VMware, Microsoft, and NetApp technologies for advanced, application-aware data protection.
- The NetApp Secure Multi-Tenancy (SMT) solution with MultiStore capability offers end-to-end data security, nondisruptive data mobility, load balancing across storage controllers, and better manageability in a multi-tenant cloud environment.

6 ACKNOWLEDGEMENTS

The authors would like to thank the following people for their contribution to the design, validation, and creation of this solution guide:

Cisco: Mike Mankovsky

NetApp: Robert Quimbey, Vaughn Stewart, John Parker, Abrar Y, Sourav Chakraborthy, Amrita Das, Niraj Jaiswal, Umesh Venkatesh, Chris Gebhardt, Jack McLeod, Mike Zimmerman, and Ryan Rothert

VMware: Wen Yu, Michael White

7 REFERENCES

TR-3749: NetApp and VMware vSphere Storage Best Practices
http://media.netapp.com/documents/tr-3749.pdf

TR-3671: VMware vCenter Site Recovery Manager in a NetApp Environment
http://www.netapp.com/us/library/technical-reports/tr-3671.html

TR-3845: “SnapManager 6.0 for Microsoft Exchange Best Practices Guide
http://media.netapp.com/documents/tr-3845.pdf

TR-3715: “SnapManager for Microsoft Office SharePoint Server: Backup and Recovery Guide
http://media.netapp.com/documents/tr-3715.pdf

TR-3505: “NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
http://www.netapp.com/us/library/technical-reports/tr-3505.html

TR-3737: “SnapManager 2.0 for Virtual Infrastructure Best Practices
http://www.netapp.com/us/library/technical-reports/tr-3737.html

VMware vCenter Site Recovery Manager Administration Guide
http://www.vmware.com/files/pdf/vcenter-srm-evaluators-guide.pdf

VMware vCenter Site Recovery Manager Best Practices for Performance
http://www.vmware.com/files/pdf/VMware-vCenter-SRM-WP-EN.pdf

VMware vSphere 4: Exchange Server on NFS, iSCSI, and Fibre Channel
http://www.vmware.com/files/pdf/vsphere_perf_exchange-storage-protocols.pdf

VMware ESX Configuration Guide
http://vmware.com/pdf/vsphere4/r40/vsp_40_esx_server_config.pdf

VMware Fibre Channel SAN Configuration Guide
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_san_cfg.pdf

VMware iSCSI SAN Configuration Guide
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf

Cisco Nexus 5000 Series Switch CLI Software Configuration Guide
http://www-tss.cisco.com/eservice/vho/lan_switch/nexus5000/mod4/docs/Nexus5000-CLI-ConfigurationGuide.pdf

Microsoft KB-888794: “Considerations when hosting Active Directory domain controller in virtual hosting environments”

Microsoft KB-875495: “How to detect and recover from a USN rollback in Windows Server 2003”

Microsoft KB-255504: “Using Ntdsutil.exe to transfer or seize FSMO roles to a domain controller”

8 FEEDBACK

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

8.1 APPENDIX: VCENTER RIGHTS REQUIRED FOR THE SNAPDRIVE ACCOUNT

To provide more security with respect to the account that is used by SnapDrive in the guest OS, create an account in Active Directory. This account does not need to be in the domain administrator group or in any local administrator group. Then create a new administrative role in the VMware vCenter server or servers at both the protected and recovery sites. Assign the following three rights to the new role and assign the role to the new user.

- **Host > Configuration > Storage Partition Configuration.** Required for performing operations such as `RescanAllHBA` on the ESX server
- **Virtual Machine > Configuration > Raw Device.** Required for adding and deleting RDM disks
- **Virtual Machine > Configuration > Change Resource.** Required intermediately when resources assigned to a VM change during disk operations

REVISION HISTORY

Date	Author	Comment
August 2010	Larry Touchette Abhinav Joshi	First release
May 2011	Sitakanta Chaudhury Suresh Vundru	Included scenarios of MS Apps (SharePoint and SQL Server) on VMDKs

Go further, faster®

www.netapp.com