BEST PRACTICES GUIDE

VMware Site Recovery Manager and Nimble Storage

Wen Yu, Nimble Storage
Ken Werneburg, VMware

Document Revision

.

Date	Revision	Description
3/8/2013	1.0	Co-authored/branded by VMware and Nimble

Table of Contents

Introduction

About Nimble Storage

Nimble Storage is the leader in flash-optimized hybrid storage solutions, providing customers with scalable, efficient, and high performance storage for all mainstream applications, user data, and server and desktop virtualization workloads. Based on the Cache Accelerated Sequential Layout (CASL™) architecture, Nimble Storage accelerates applications, protects more data, and empowers IT to take on new projects and drive growth.

About Site Recovery Manager (SRM)

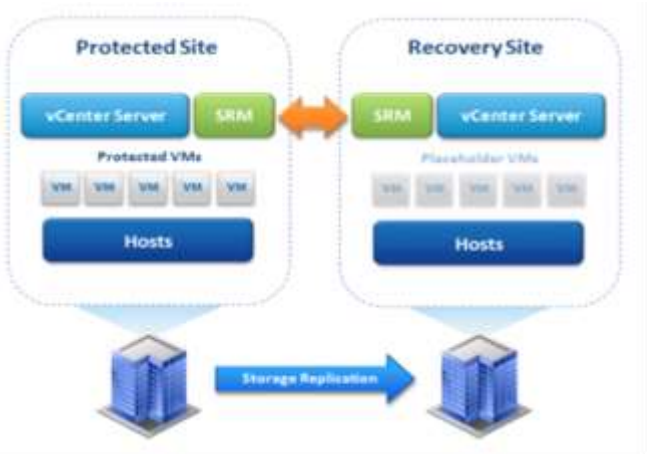

Simplifies, automates DR

- Setup
- Testing
- Failover
- Auditing

VMware® vCenter™ Site Recovery Manager (SRM™) is an extension to VMware vCenter that provides disaster recovery capabilities to VMware customers. Site Recovery Manager enables simplified automation of disaster recovery. For more information about SRM, please visit the following VMware resource page (http://www.vmware.com/products/site-recovery-manager/overview.html).

Relevant Terminology

SRM Terminology

Terminology	Definition
Sites	Protected Site and Recovery Site
Placeholder VM	Protected VM that shows up as a placeholder in the Recovery Site vCenter Server inventory
Placeholder Datastore	Datastore to host shadow VM config (.vmx) file
Array Managers	Allows SRM to communicate with Nimble Array
SRA	Storage Replication Adapters
Consistency Group	All VMs in a consistency group replicate and recovery together
Protection Group	Groups of VMs that use the same datastore consistency group
Recovery Plan	Steps required to test, planned migration or disaster reclamation
Reprotect	Configures reverse replication, in preparation for failback

Nimble Storage Terminology

Terminology	Definition
Zero Copy Clone	Space efficient volume cloning technology built-in with CASL
Replication Partners	Pair of Nimble arrays configured for replication
Volume Collection	Collection of Nimble volumes that have a common snapshot/replication schedule
Volume Collection Synchronization	Method of synchronization to quiesce I/O (either with Microsoft VSS/Oracle/VMware vCenter) during snapshot creation
NPM	Nimble Protection Manager (package that needs to be installed for VMs with RDM & in-guest attached storage, to achieve application consistent snapshot)

SRM and Nimble Storage Terminology Mapping

VMware SRM	Nimble Storage
Consistency Group	Volume Collection
Local Device	Local volume
Remote Device	Replica volume
Datastore	Volume provisioned to ESX cluster
Array Pair	Local + Remote Nimble Array

About Nimble Storage Integration with SRM

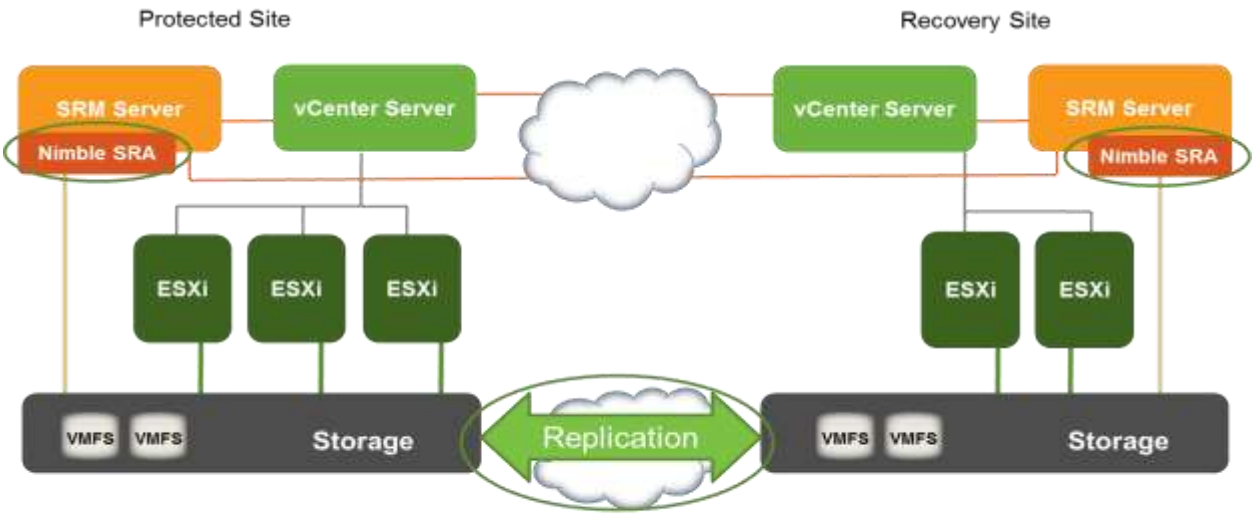

Nimble Storage fully integrates with VMware SRM to automate DR protection, test and recovery for the storage layer. Major points of integration between VMware vCenter Site Recovery Manager and the Nimble Storage Arrays are done via a “Storage Replication Adapater” (SRA) written by Nimble to the specifications provided by VMware. The SRA allows for a number of storage interaction workflows to be initiated from Site Recovery Manager, such as: Discovery; Test Failover; Cleanup; Planned Migration; Disaster Recovery; Reprotect.

Discovery: SRA helps SRM discover Virtual Machine File System (VMFS) datastores that are configured with cross-site replication. Arrays are pre-configured with replicated devices or consistency groups, which are then presented to the vSphere clusters. In the SRM server the SRA must be installed to allow visibility to these replicated devices. After the SRA is installed and configured, replicated devices are represented within SRM, available for use in workflows.

Test Failover: Creation of test copies of data at the recovery site allows SRM to execute DR tests without interrupting production systems or interfering with ongoing replication of the primary data sources. The "Test Failover" function of the SRA creates a zero copy clone of the replica volumes associated with the Protection Groups in a recovery plan, and presents the corresponding replicated VMFS datastore(s) to the recovery cluster.

Cleanup: The SRA will remove the clone created when DR testing has been completed. NOTE: Proper volume unmounts and removal tasks will be performed by vCenter prior to removal of replica clone volumes from the Nimble array side.

Recovery (Planned Migration): A planned migration will perform an initial synchronization, shut down the protected site virtual machines, and then ask the SRA to set the appropriate Protected Site datastore group volumes to "Read Only". It will then take another quiescent up-to-date snapshot, and replicate this latest snapshot to the Recovery Site. If there are any failures while executing these tasks, a planned migration workflow will halt the recovery and allow administrators to fix any problems before proceeding.

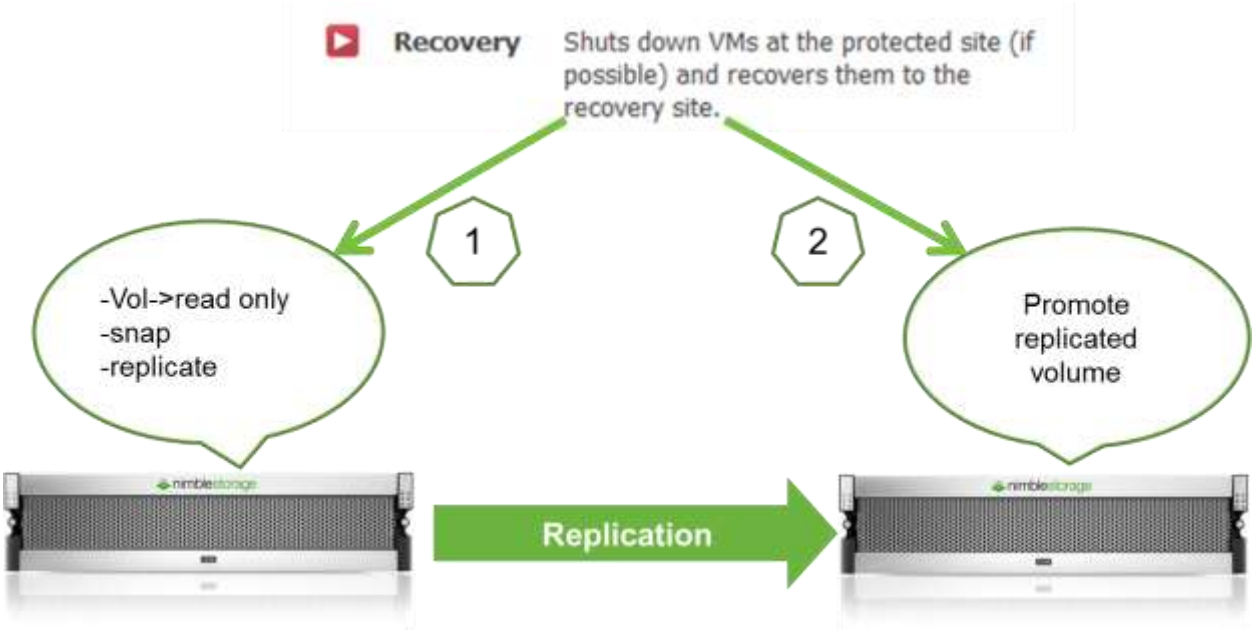

Recovery (Disaster Recovery): Very similar steps are run during a "Disaster Recovery" workflow as to those run during a "Planned Migration". With a DR scenario a rapid recovery time is more desirable than a lack of errors, so instead of stopping on errors a DR workflow will carry on running irrespective of the results from the SRA commands. The SRA promotes the latest replicated replica volume even if synchronization and VM shutdown cannot complete.

Reprotect: During a failover, the SRA will demote all volumes on the datastore group that had undergone a recovery on the Protected Site. Once failed over there may be a requirement to protect the virtual machines back to the initial site. When running a "Reprotect" workflow in SRM, the SRA invokes a snapshot of the volumes in the recovered datastore group, then reverses replication to the original Protected Site

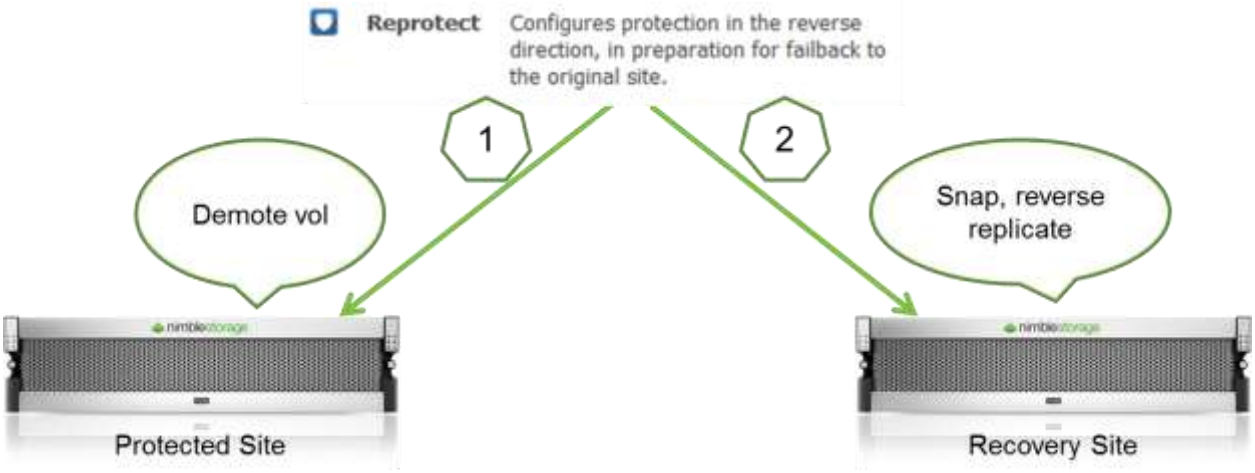

DR Consideration and Best Practices

Below is a list of core areas of considerations when designing a workable, repeatable and auditable disaster recovery strategy:

Compatibility

First and foremost, make sure the solution you are about to deploy is listed as supported by VMware and Nimble Storage. Nimble Storage Replication Adapter (SRA) is fully certified and supported by VMware. Follow the two links below to ensure you have a supported combination of hardware/software:
VMware SRM HCL:
http://www.vmware.com/support/srm/srm-compat-matrix-5-1.html
Nimble Storage HCL:
http://support.nimblestorage.com/download/documentation/sra/1-4-2-1/Nimble_SRA_5.0_5.1_Release_Notes_PN990-0005-004.pdf

Application Consistency

This is one of the most important considerations for disaster recovery protection. When a backup copy of the application data is being created, it is a best practice to ensure I/O is fully

quiescent for the virtual machines being protected. Doing so ensures recoverability of application data from a point-in-time backup copy.
Nimble Storage has built-in integration with both Microsoft Volume Shadow Service (VSS) framework and VMware Tools VSS implementation. The type of disk connectivity method dictates the volume collection synchronization method. The three virtual storage connectivity methods are discussed below, along with the design consideration for each:

NOTE:

- Ensure only VMs needing DR protection are placed in volumes configured with replication. VMs that that need to failover together should also reside on the same volume collection, to ensure consistent snapshot and replication schedule

- A manual volume snapshot initiated from Nimble UI or vCenter Server plugin would NOT trigger the VSS requestor to request for application quiescing. Ensure a snapshot schedule is in place for the Volume Collection for proper VSS application quiescing

- Both vCenter Server and SRM databases should have a regular snapshot and replication schedule. It is also recommended to add a manual schedule after each DR test, recovery, and reprotect task. Remember a scheduled snapshot is needed to trigger proper VSS quiescing of I/O

Here's a quick summary of each virtual storage connectivity method, along with design considerations:

Connectivity Method	Volume Collection Synchronization	NPM Installation	Snapshot Taken	Design Considerations
Raw Device Mapping(RDM)	Microsoft VSS	Inside the guest OS (NPM requestor & provider used for snapshot)	Nimble Array based snapshot	• No VMware snapshot taken
VMDK	vCenter	None (VMware VSS requestor/provider used for snapshot)	VMware snapshot + Nimble Array based snapshot	• No Log truncation for MS Exchange • Avoid too many VMs in the same datastore
In-guest Attached iSCSI	Microsoft VSS	Inside the guest OS (NPM requestor & provider used for snapshot)	Nimble Array based snapshot	• Manual scripting work needed to mount in-guest attached iSCSI storage during test recovery, and recovery

RDM(Physical Compatibility Mode) for application disks, VMDK for OS disk

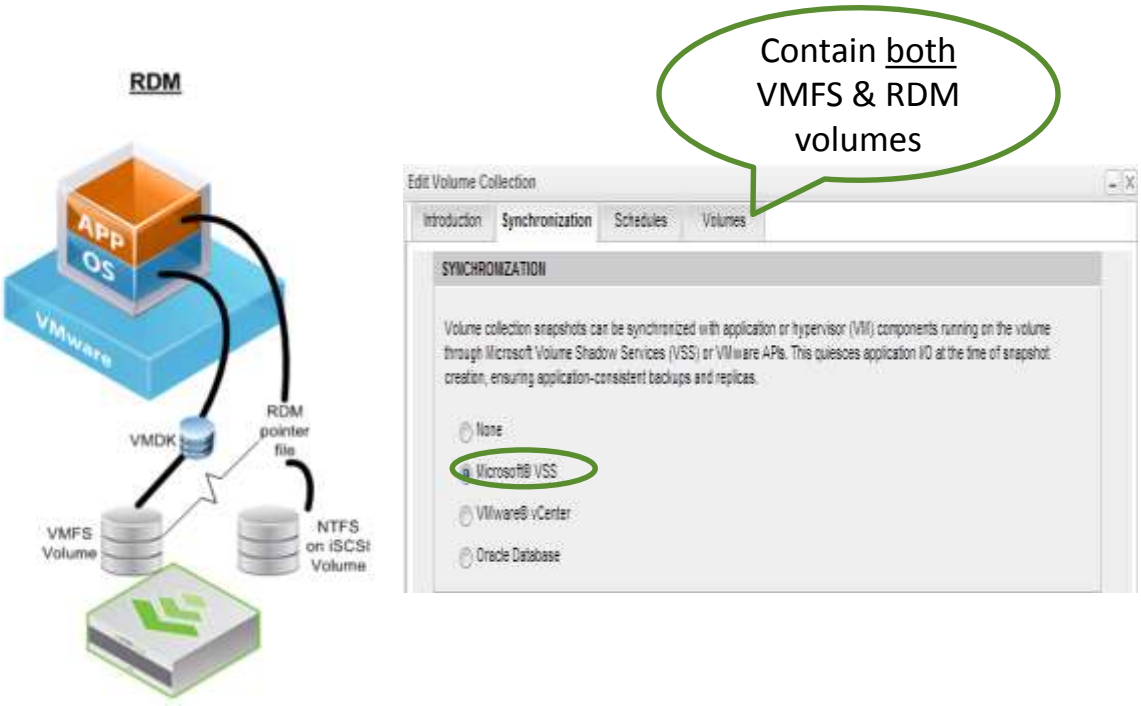

This method involves using RDM for application data and VMDK for VM guest OS. With this method, VMware ESXi server will simply ignore the RDM disk during VMware snapshot operation, meaning vSphere will only create a VMware snapshot for the O/S VMDK. As for the application disk that is running as RDM, the Nimble VSS hardware provider will be used for snapshots. Therefore, it is imperative to ensure the Volume Collection containing the RDM and VMDK volumes has “Microsoft VSS” synchronization selected.

NOTE:

- With “Microsoft VSS” synchronization, there would be NO VMware snapshot taken by ESXi servers. The Nimble VSS hardware provider will be leverage for taking the snapshot, after the VSS writer has successfully freeze incoming I/O requests

- The Nimble Windows Toolkit (Nimble Protection Manager/NMP) needs to be installed on the VM that has RDM storage attached

VMDK for both application and OS disks

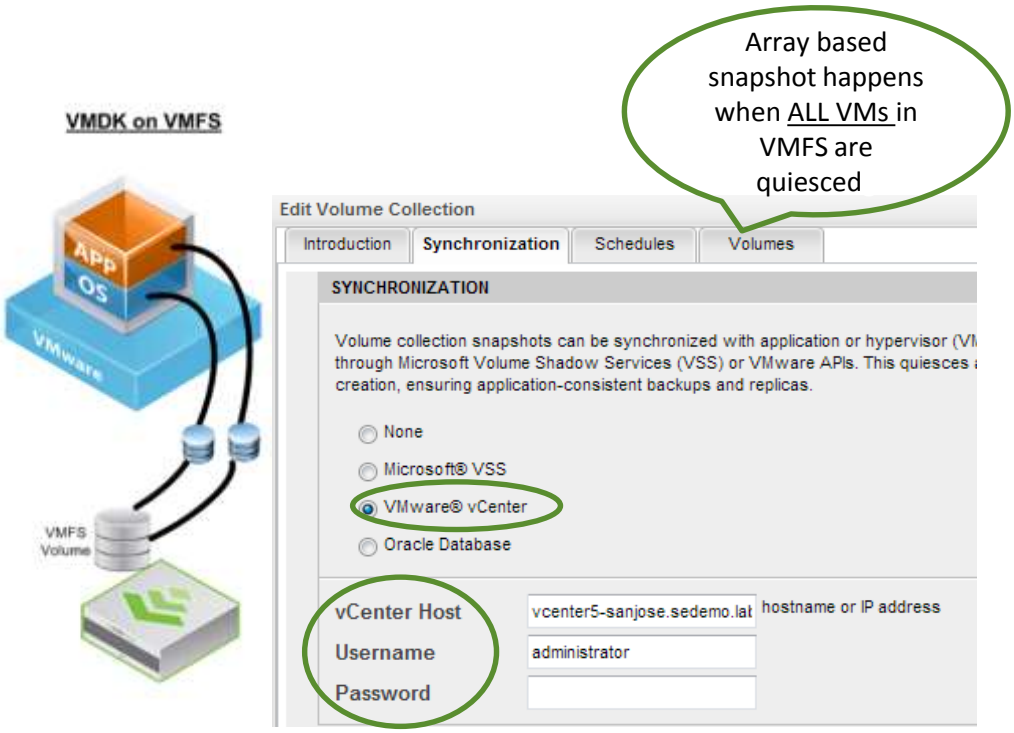

This method involves using VMDK for both VM guest OS and application data. With this method, keep in mind that ALL VMs in a given VMFS volume need to be quiesced to take VMware snapshot, follow by the Nimble array taking an array level snapshot. It is a best practice, if using this method, to limit the number of virtual machines that reside on the given VMFS volume on which the quiescence is being issued. This will minimize both the number of VMs impacted by the suspend operation as well as improve the speed with which the snapshot can be taken.

 NOTE:

- The current VMware implementation of the software VSS provider does NOT truncate logs for Microsoft applications such as Exchange, SQL Server or Sharepoint. If you have an integrated backup application such as Commvault that could be invoked to truncate the logs, be sure to leverage that. If not, you could:
    - enable circular logging in Exchange;
    - consider in-guest/RDM mounted storage, and
    - build custom script to invoke during backup to truncate the Exchange logs.

In-guest attached iSCSI storage for Application Data, VMDK for OS disk

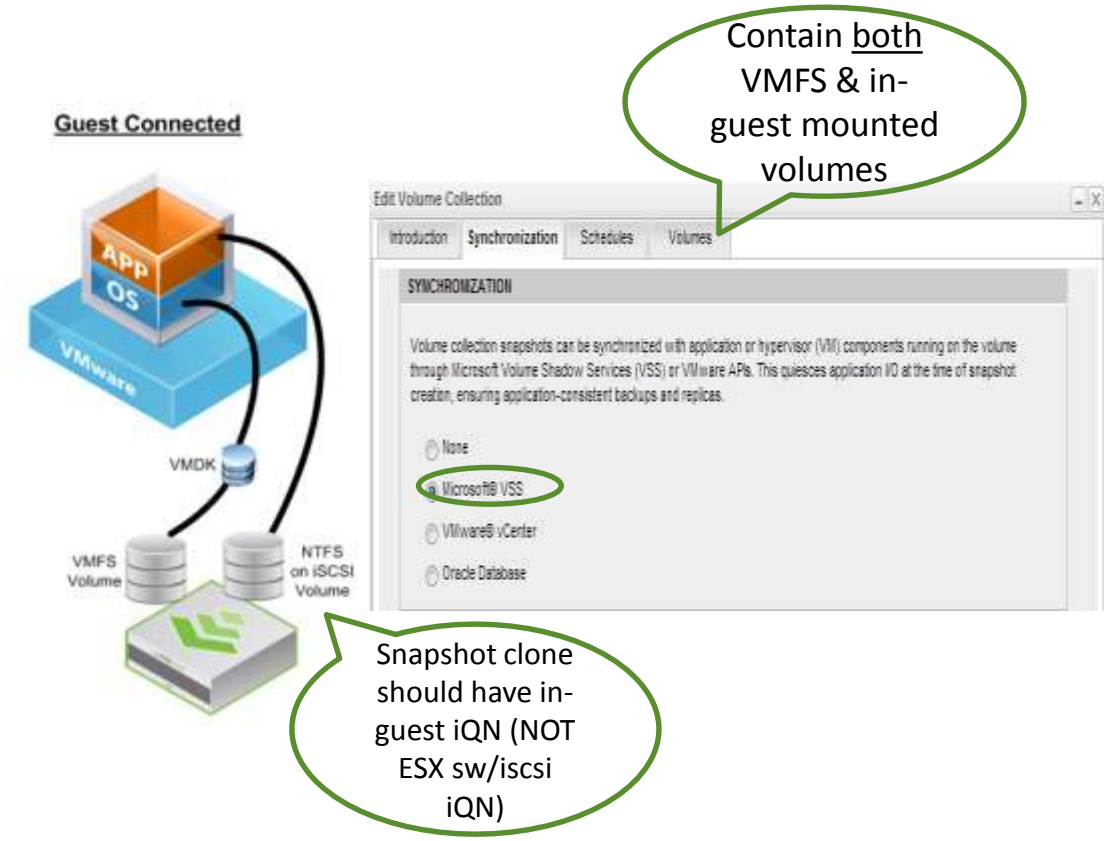

This method involves using a VMDK for the VM guest OS disk, and in-guest attached iSCSI storage for application data. With this method, the in-guest attached storage will bypass the ESXi storage stack, and appear as network traffic to ESXi server. Typical use cases for this configuration are:

- Using Microsoft Clustering Service (MSCS) on iSCSI protocol, and
- Reaching beyond the 2TB VMDK size limitation.

Just like the other two connectivity methods, there are design considerations and tradeoffs.

NOTE:

- SRM does not know about in-guest attached storage, therefore during recovery extra steps are required to mount these volumes for each VM that uses this type of storage.

- The Nimble Protection Manager (NPM) needs to be installed on the VM with in-guest attached iSCSI storage

With the correct Volume Collection Synchronization setting, each scheduled snapshot is application consistent. Therefore, when the scheduled snapshots get replicated to the Nimble replication partner array, the new replica is also application consistent.

RTO and RPO

RTO (Recovery Time Objective) and RPO (Recovery Point Objective) are two key factors in architecting a disaster recovery solution that meets the Service Level Objective (SLO) for the business.

RTO refers to how quickly a business can recover from the disaster, or specifically how long it takes to execute the recovery process to bring back business services. RTO is typically measured in minutes or hours. Site Recovery Manager is what automates the recovery and maximizes the RTO of a vSphere DR solution.

RPO refers to how far back in time the data in the environment will be, once the business service is restored. RPO is typically measured in minutes. Nimble Storage replication provides the aggressive RPOs demanded of a DR solution.
The integrated SRM solution with Nimble Storage will enable end users accelerate their RTO, as the configuration, DR test, recovery and failback are simplified and automated. The best practice is to conduct regular DR testing to ensure the RTO is met consistently and repeatedly.

For RPO, here are the key design considerations for each factor that affects RPO:

RPO Dependency	Considerations
Size of dataset between scheduled snapshots	Gather statistics on the snapshot storage capacity utilization (obtain average size and maximum for each volume collection that hosts critical application data)
WAN link connection type/speed	Is storage replication the only consumer of WAN bandwidth? Factor in other usage of the WAN connection bandwidth
Scheduled replication frequency	Snapshot frequency is derived by calculating the amount of time replication takes for average and max dataset size for snapshot. The frequency should be >= replication time for average/max dataset

NOTE:

Both VMware vSphere and Nimble Storage have built-in features to reduce replication footprint, to make more efficient use of WAN bandwidth.

VM Swap File Location in vSphere:
Place all the virtual machine swap files in a dedicated VMFS datastore without replication enabled. By doing this it will eliminate replication of transient virtual machine data, reducing the amount of bandwidth needed to protect a virtual machine and saving on time and cost of replication. VM swap files are recreated automatically by SRM when the recovery copies of the virtual machines are powered on. This setting can be found in the Cluster Features setting in vCenter Server:

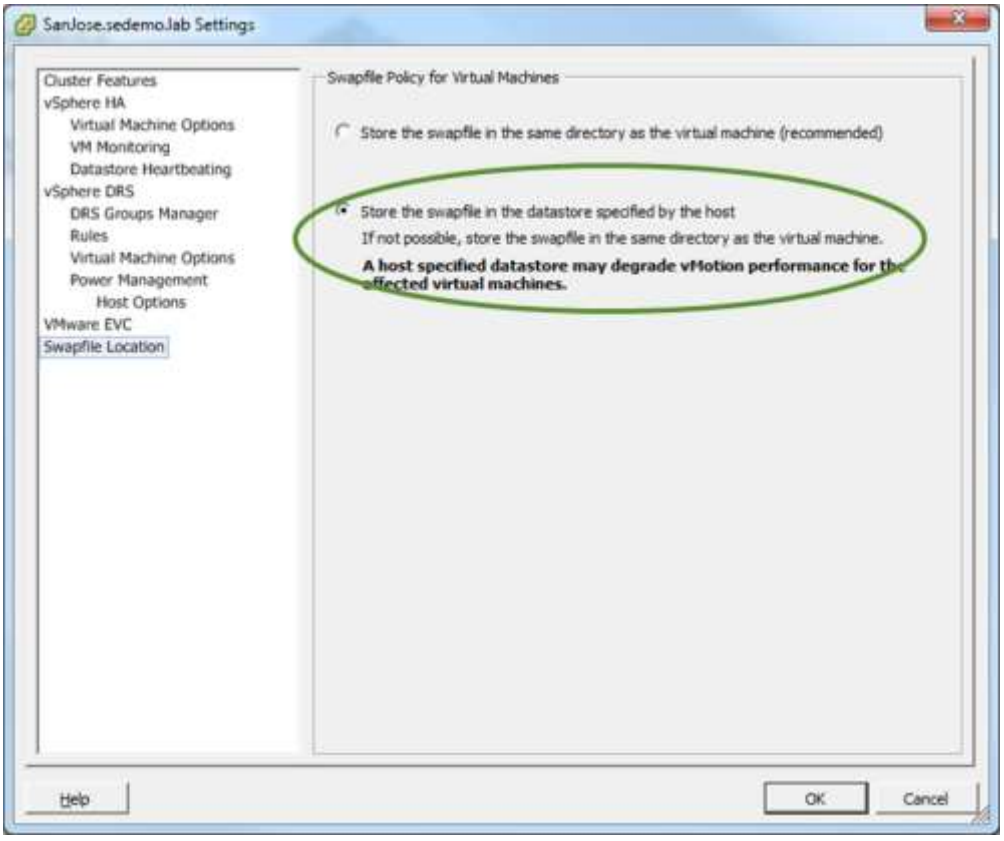

Replication Policy:
Replication data reduction is enabled by default. All volumes created in Nimble array have "Compression" enabled. Leaving it enabled is a best practice as the snapshot data is also compressed. Additionally, the Nimble replication engine eliminates data duplication by only replicating the delta between the primary and replica volume. By sending only deduplicated and compressed data, the amount of time necessary to replicate a virtual machine is shrunk, enabling better RPOs for the environment.

If the WAN connection is shared with other critical applications, QoS policies could be created to limit the amount of bandwidth used by Nimble Storage at the network layer (switch/firewall) or at Nimble Storage Replication Bandwidth configuration menu:

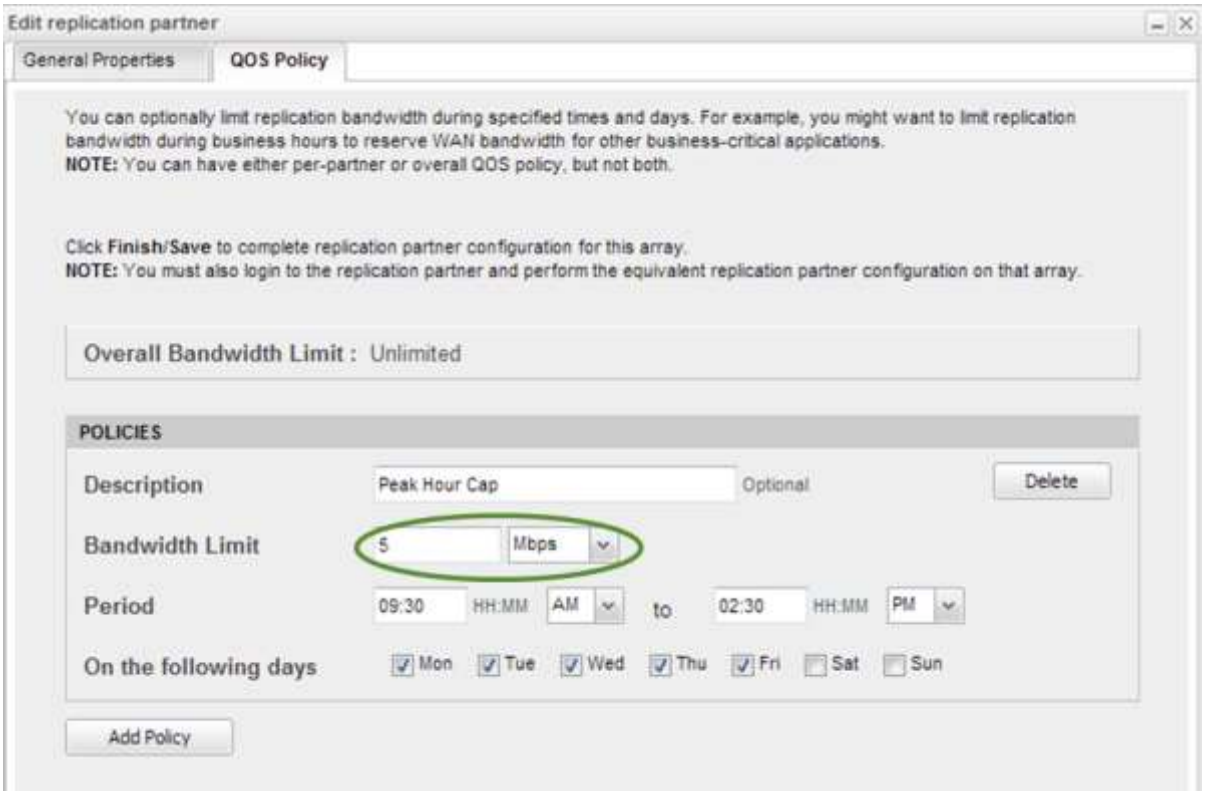

Base Infrastructure Service

Base infrastructure service such as Active Directory (AD), Domain Name Resolution (DNS), and HTTP are core to the business services applications. Therefore, it is imperative to ensure these services are running at all times in Recovery Site to ensure successful DR testing, recovery and failback.

NOTE:

AD should not be replicated using storage replication technology, as it could potentially cause a USN rollback scenario. Refer to Microsoft KB 888794 for further details: (http://support.microsoft.com/kb/888794)

Configure high availability for AD by having at least two instances of AD server, both configured as a Global Catalog Server. Additional considerations for AD are listed in the Test, Recovery, and Reprotect Section.

In almost all scenarios it is most favorable to have fully active Active Directory and DNS services already running at the recovery site. This will minimize the amount of change necessary to fail over an environment to the recovery location, granting a better recovery time. It will also ensure no last-minute changes are needed for the authentication/DNS environments that may prove difficult to overcome during a disaster, giving you more reliable DR.

Network

If both your datacenters are located on the same Layer Two network or on a stretched VLAN, you could skip this section. If you have to change IP addresses for a small number of virtual machines during DR test, recovery, you could leverage the VM Recovery Properties wizard in the Recovery Plan, for each virtual machine:

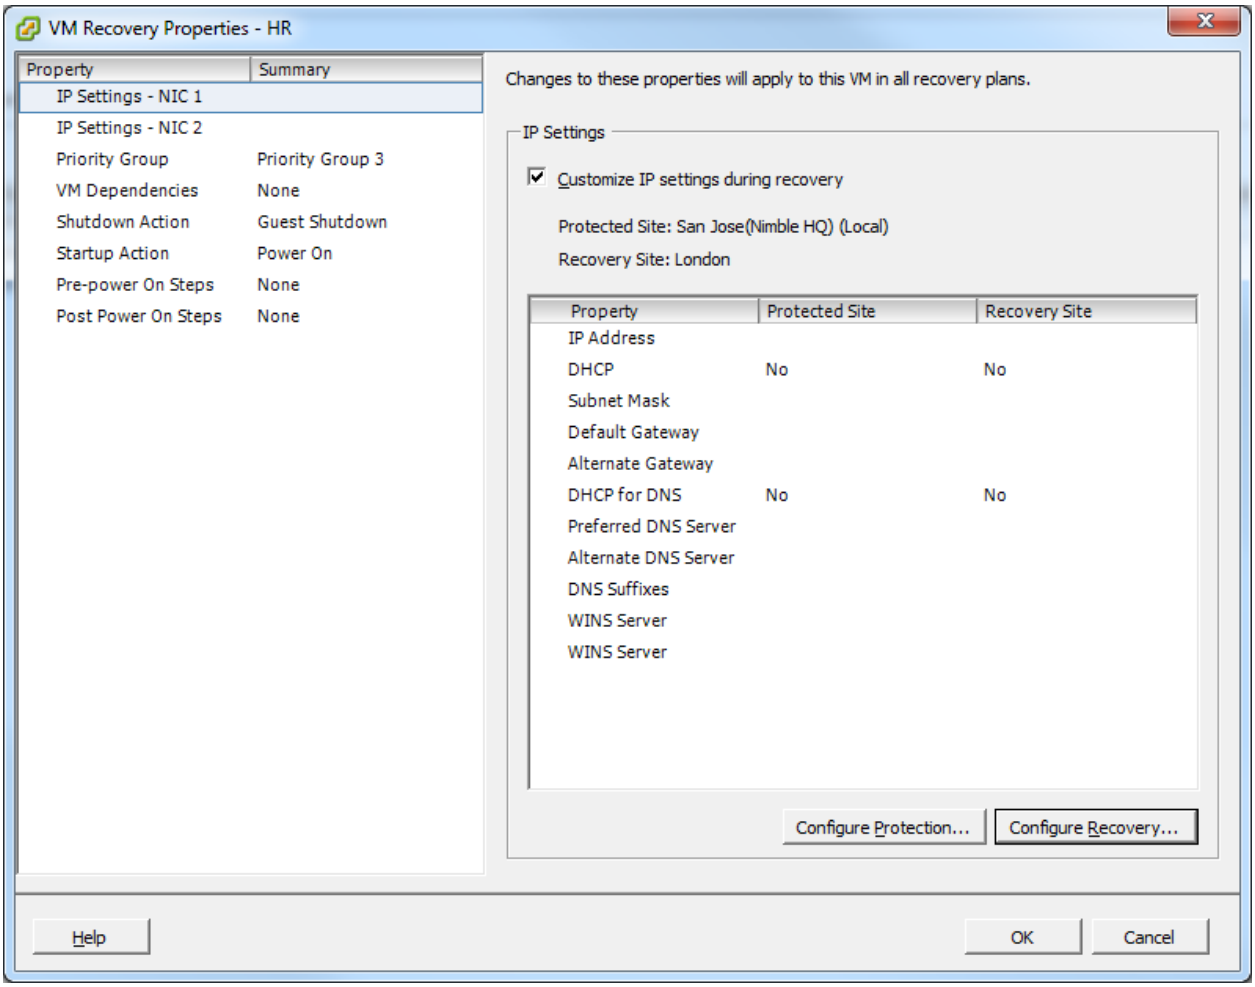

NOTE:

Ensure the Network Mapping is configured properly for the Recovery Site:

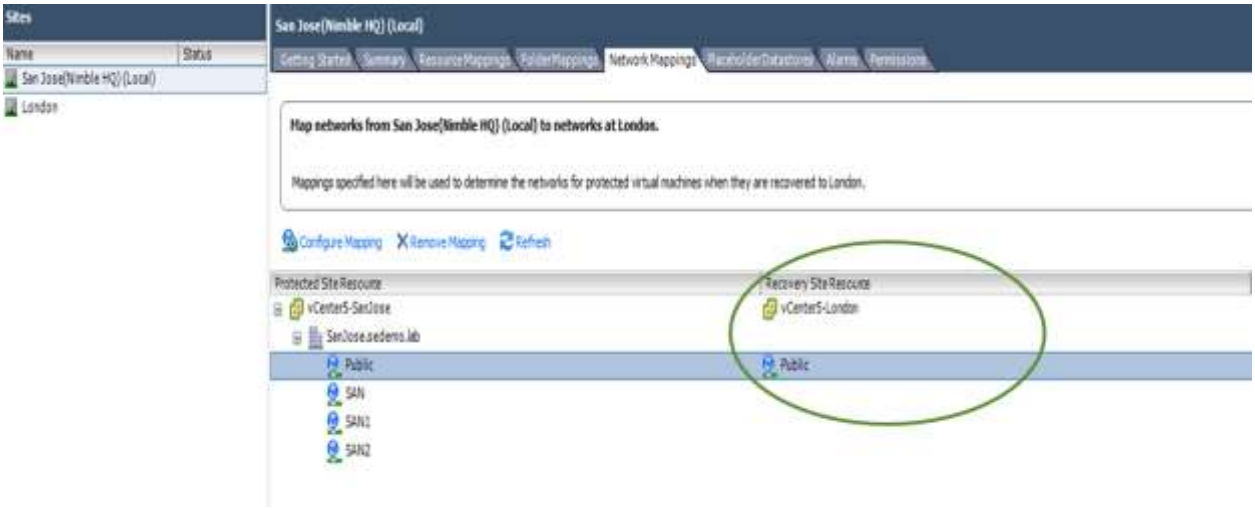

If you have to change IP addresses for a large number of VMs (for example, hundreds of VMs), then you should consider automating this process with the SRM DR IP Customizer tool. More information about this tool can be found here: http://pubs.vmware.com/srm-51/index.jsp?topic=%2Fcom.vmware.srm.admin.doc%2FGUID-9794F585-C168-48D6-8866-E8519E768278.html

If it is possible to run the recovery site with the same network IP ranges, this will allow you to fail over virtual machines without reconfiguration. For optimal recovery time stretched layer 2 networking is recommended to avoid the time required to change networking for all virtual machines.
If a stretched layer 2 network is not achievable in your environment, SRM can automatically reconfigure networking for virtual machines upon failover as outlined above.

DR Testing

One of key aspects of meeting DR RPO and RTO is testing. It is a good practice to test DR recovery process on a regular basis. Here are the key design considerations for DR testing with SRM and Nimble Storage:

 NOTE:

Ensure base infrastructure service could be verified during DR testing. This means all critical application VMs could join the domain, end users and applications could authenticate against AD, and the application components could communicate on the network. By default, SRM has a test bubble network that gets created without real network connectivity. Follow the recommendations and best practices below to truly verify infrastructure service availability during DR test:

Base Infrastructure Service	General Recommendation	VMware Best Practice	Nimble Best Practice
Active Director (includes DNS/HTTP services)	• Ensure AD servers are configured as Global Catalog Servers • Clone AD from Recovery Site during DR test	N/A	• Configure dedicated volume to host base infrastructure VMs; during DR testing, create a zero copy clone of the volume • Destroy the zero copy clone after DR testing • The same volume could be used as the Placeholder datastore
Network	• Configure dedicated private VLAN/network for DR testing purpose	• If vMotion network is configured in the Recovery Site, leverage this port group for simplicity • Ensure the protected VMs are configured with appropriate address for the dedicated test network	N/A

NOTE:

For applications that require the FSMO roles in the Active Directory forest. Refer to Microsoft KB 255504 (http://support.microsoft.com/kb/255504) for procedure to seize the FSMO roles during DR test.

Here's a high level workflow of infrastructure preparation and cleanup for DR testing:

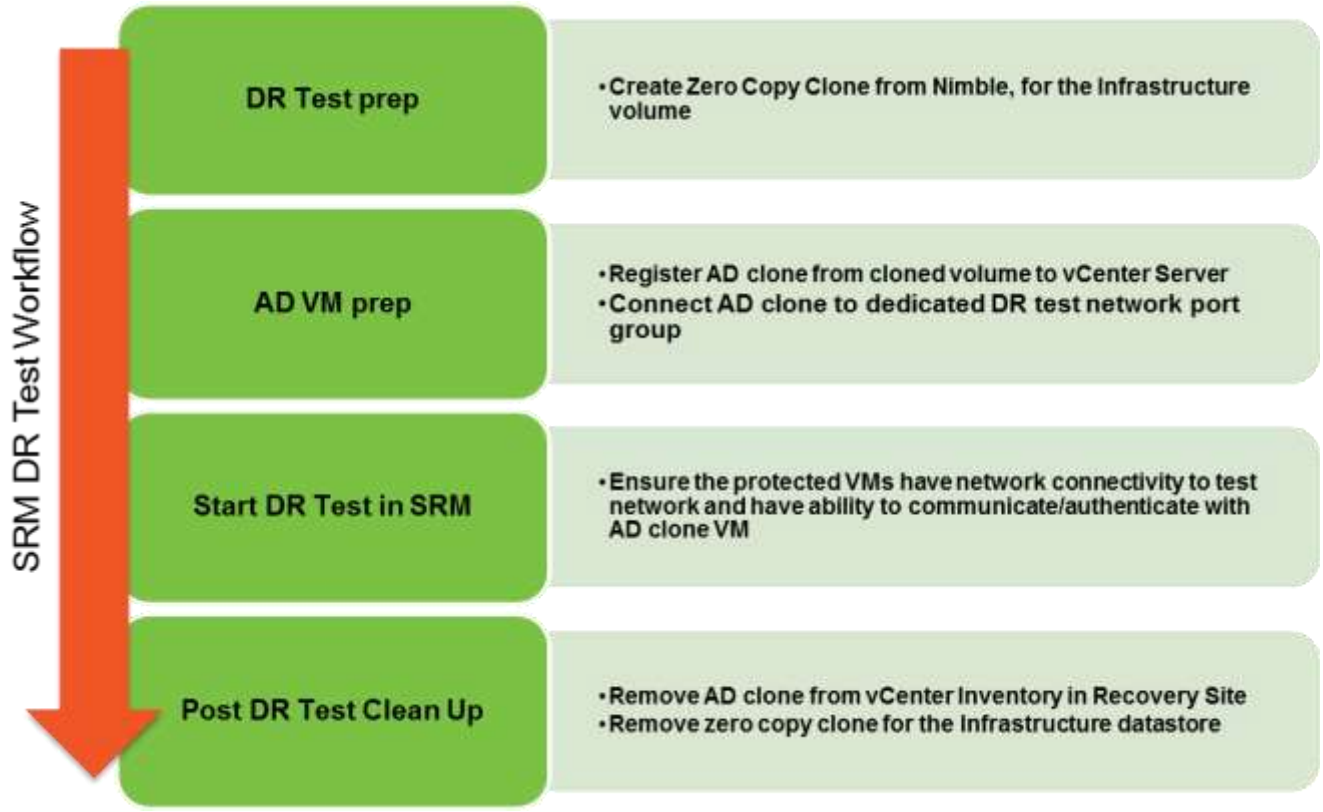

Audit and Reporting

SRM has built-in audit and reporting functionality to log DR test, failover (planned migration and DR), and reprotect activities. It is imperative to protect the databases backing these logged activities, and cross replicate the volume collections between the Protected Site and Recovery Site.

Reference Materials

Nimble Storage and VMware vSphere Best Practice Guide

http://info.nimblestorage.com/bpg-vsphere-5.html

Microsoft SQL Server Best Practice Guide

http://info.nimblestorage.com/bpg_sql-server.html

Microsoft Exchange 2010 Best Practice Guide

http://info.nimblestorage.com/bpg_exchange.html

Architecting Storage in Virtualized Environments

http://info.nimblestorage.com/storage-for-virtualization-1.html

Nimble Storage, Inc.

2740 Zanker Road., San Jose, CA 95134

Tel: 877-364-6253; 408-432-9600 | www.nimblestorage.com | info@nimblestorage.com