BEST PRACTICES GUIDE

Nimble Storage Best Practices for Microsoft Hyper-V R2 and R3

Table of Contents

Hyper-V Availability Reference Architecture

Hyper-V High Availability and Disaster Recovery solutions can be implemented in many different ways. This reference architecture weighs the pros and cons of individual options for Hyper-V implementation to provide the most feature-rich protection solution using clustered Windows 2008 R2 Hyper-V, Windows 2012 Hyper-V (R3) and Nimble Storage.

The following are some of the primary solution benefits provided by these best practices:

- **Support for Hyper-V Live Migration**: Microsoft Hyper-V requires Failover Clustering to perform Hyper-V Live Migration between host servers. This greatly reduces the amount of effort required to migrate servers due to maintenance, load balancing, or hardware consolidation.
- **Performance Policies**: Traditional storage devices force application writes into static-sized block or page containers that do not optimize storage space. Nimble Storage developed the patent-pending CASL file system which uses variable-length blocks that precisely match application write sizes to maximize storage space. Variable-length blocks combined with real-time inline compression greatly reduces the footprint for data storage, snapshots, and replication, which allows you to store more data and greatly reduce bandwidth costs especially over Wide Area Networks.
- **SCSI Unmap:** Provides reclamation of storage space on thin provisioned volumes when data is deleted.
- **Application Awareness**: Storage replication by itself can put your data at risk for applications that perform transactional write processes, such as databases. Nimble Storage provides application integration to ensure that they properly flush their write buffers to a quiescent state prior to triggering point-in-time operations like snapshot backup and replication.
- **Zero-Copy Cloning**: Nimble Storage greatly reduces the amount of storage required for virtual machines by eliminating duplicate files common to multiple operating system images.
- **High Availability**: When a system failure occurs, Microsoft Failover Clustering quickly restarts virtual machines on surviving hosts automatically. This reduces the amount of effort required to manually perform virtual server recovery. Nimble Storage fully supports Microsoft Failover Cluster and Hyper-V technologies, providing high-speed fault-tolerant storage.
- **Off-site Disaster Recovery**: Recovering production applications to an off-site disaster recovery location using Nimble Storage WAN-efficient replication provides you with fast recovery and business continuity cost-effectively in the event of a catastrophic site outage.

The following best practices will guide you through implementation and management of this architecture to maximize your Hyper-V system availability and off-site recovery with minimal effort.

Clustered Hyper-V Server High Availability Architecture

The primary drawback of server virtualization is that system outages can now affect multiple machines simultaneously. To protect against such impact, you should implement Microsoft Hyper-V using the Windows Server Failover Cluster Role which provides automatic recovery in the event of a server failure. This greatly simplifies the management effort required to recover applications after an outage. This architecture also allows the use of Hyper-V Live Migration to proactively move virtual machines between host servers for better application scaling.

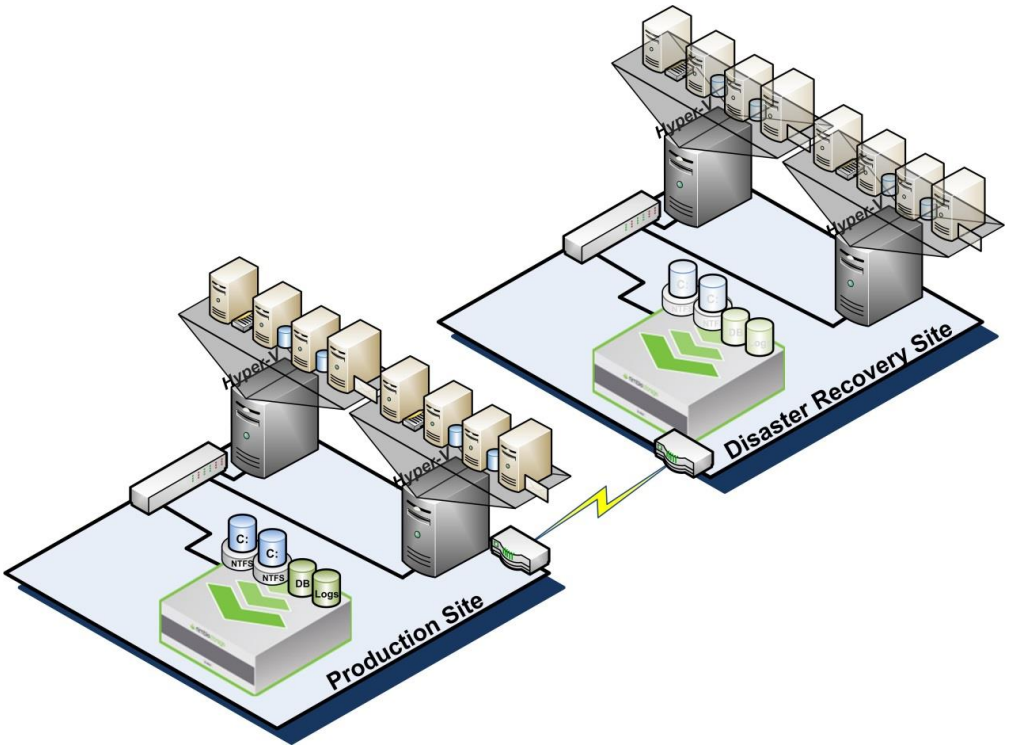

Implementing a Microsoft Failover Cluster requires shared SAN storage that all hosts of the cluster can access. Nimble Storage provides a high-performance, fault-tolerant storage platform that is fully compatible with Microsoft clustering. To ensure proper compatibility with your server platforms, follow Microsoft best practices for implementing your cluster.

Microsoft Windows Server 2012 Hyper-V R3 permits Live Migration of running virtual machines without requiring shared storage. However, this technology is used for proactively migrating virtual machines and does not provide high availability in the event that a host server fails. Shared storage – such as Nimble Storage CS arrays - are required to provide high availability within a Hyper-V clustered environment.

Hyper-V Storage Architecture

Implementing highly available Hyper-V clusters requires shared storage accessible by all hosts participating in the cluster. Nimble Storage provides a robust storage architecture that gives you fully redundant hardware and seamless access to volumes from all Hyper-V cluster nodes.

The following guidelines will help you to implement your Nimble Storage for maximum benefit for Microsoft Hyper-V.

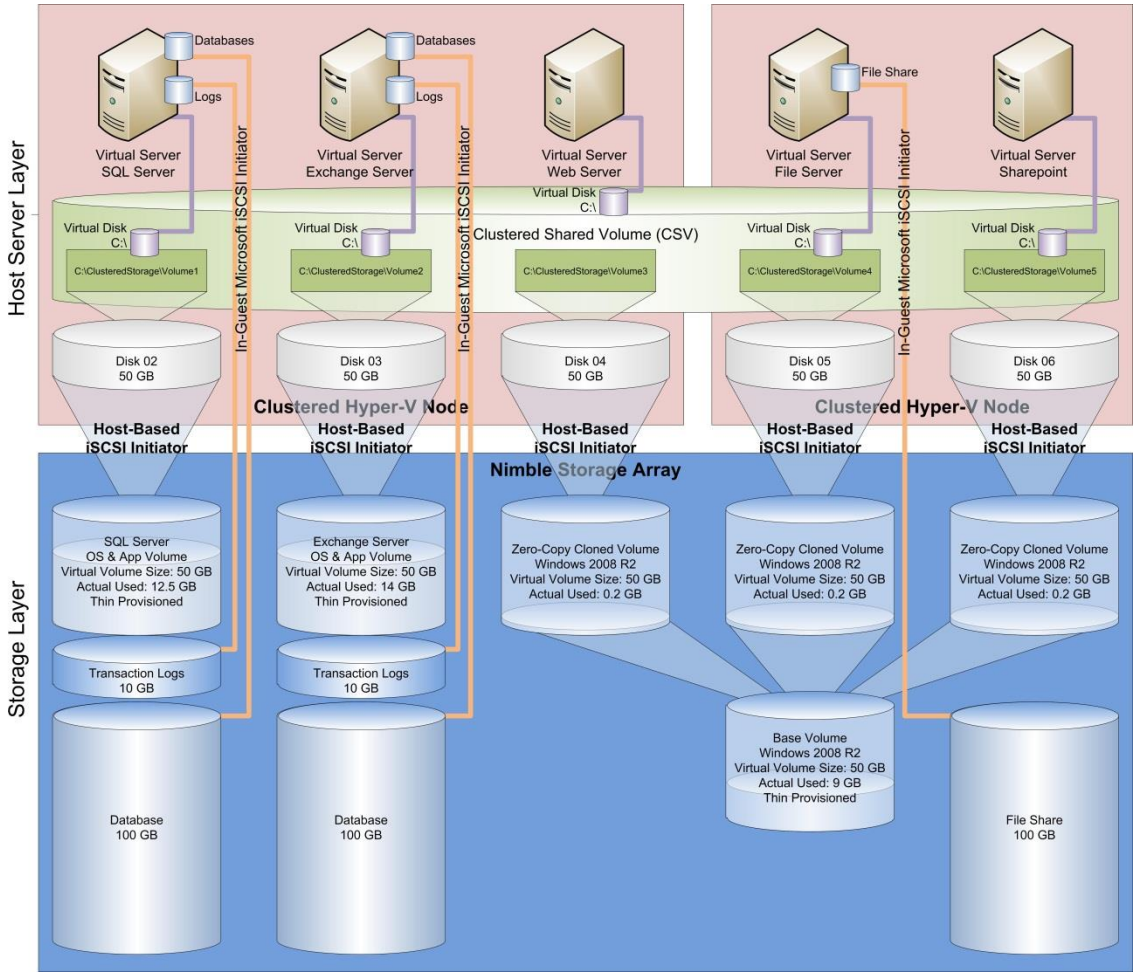

Separate Volumes for OS/Applications and Data

When creating a new virtual machine, you should separate the operating system and application binaries volume from the data volumes. Operating systems and application binaries change infrequently enough that simple volume crash consistency is acceptable therefore place OS virtual disks on Cluster Shared Volumes (CSV).

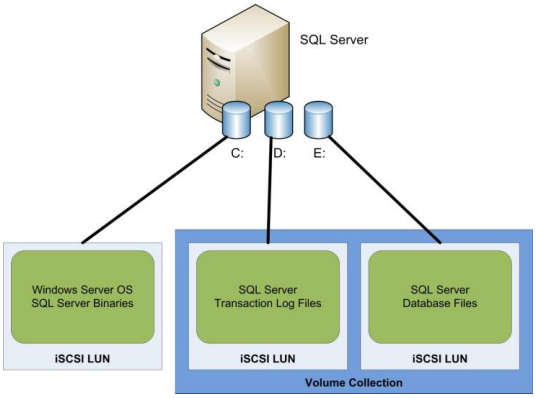

You should attach database volumes from the Nimble array to the guest virtual machine running the database application. This separates data from the operating system and application to allow cloning for development and testing, which gives you quick access to production data sets without wasting storage space to make copies of the data. Data volumes tend to change constantly and typically have more critical protection needs. For example, database applications usually write changes to a transaction log prior to writing to the database files. This allows them to recover any partial write activity in the event of a catastrophic system failure, such as a sudden power outage. If database applications did not perform this write process (WAL Algorithm) then the database could be left in a non-recoverable, and therefore non-trusted, state that forces a complete restoration from a backup. Therefore, it is important to protect both the transaction logs and database in a coordinated fashion when performing any type of backup operation. Nimble Storage arrays provide functionality that allows you to group volumes that need to be mutually consistent into the same Volume Collection.

Use Volume Collections

A volume collection allows you to schedule the frequency and retention of snapshots as well as replication to other Nimble Storage arrays. A volume collection can coordinate protection activities between separate yet related volumes (such as a database's transaction log and database file volumes) to ensure that databases are snapshot with application consistency. The volume collection integrates with Microsoft VSS, which triggers it to momentarily quiesce the write activity of the file system or application respectively to ensure data integrity of the point-in-time backup.

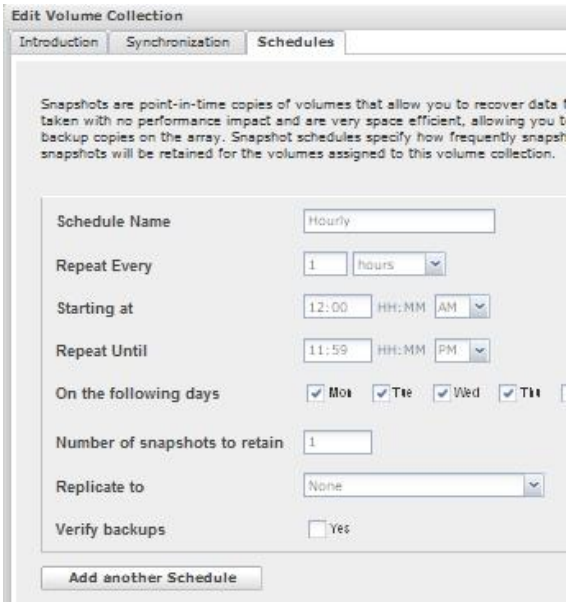

Use Guest Connected iSCSI Volumes for Storing Data

You should store data on Nimble Storage iSCSI volumes rather than on virtual hard disks. This method of data storage connectivity provides the best solution for data protection, application consistency, and off-site replication, as well as performance enhancements, by making complete use of Nimble Storage array optimization features. You should install the Nimble Windows Integration Toolkit (available for download from the Nimble Storage support web site) on each of your guest virtual machines that store data.


Do not store data in virtual hard drives.

Nimble Storage does not currently support storing application data in virtual hard drives (neither VHD nor VHDX) if the data requires a quiesce prior to snapshot. This includes database and transaction logs for applications such as SQL Server and Exchange.

Implementing Storage for Fail-over Clusters

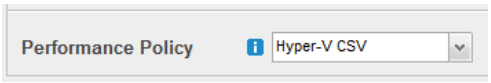

The first step to configuring shared cluster storage for Hyper-V is to provision a volume on the Nimble Storage array. Use Performance Policies that pre-configure new volumes using optimized configuration settings specific for different usage scenarios. For example, the Hyper-V CSV performance policy is tuned to use 4 KB volume blocks to provide the best performance for Windows storage LUNs. The Hyper-V CSV performance policy also includes in-line compression and high-performance caching. Thus, use the Hyper-V CSV performance policy for your cluster shared volumes.

When provisioning storage for your data, use a performance policy specific for your application. Nimble Storage provides performance policies for major applications such as Microsoft SQL Server and Exchange or create your own. For example, you might have large files that are already highly compressed, such as a video or image server, that perform better with larger block sizes and no compression. Use the Nimble Storage array's monitoring tools to view your volume performance under simulated production loads to better understand your unique application best practices.

Isolate your storage using access control lists (ACL) that include an initiator group for the cluster with each node's IQN added to the initiator group. This will reduce management and improve security by isolating volumes to only the machines that should have access to that data.

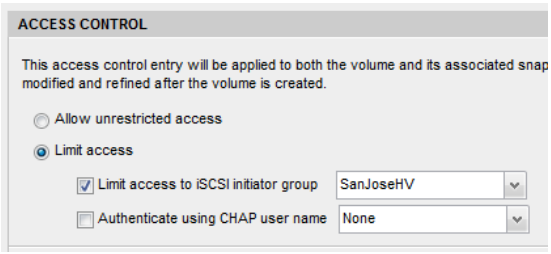

Using a server's IQN is preferable to using the IP address since they are tied to the host server and not a particular NIC. Thus a server with multiple NIC ports and IP addresses will have a single IQN. In-addition the IQN will only change if the host name changes but not if IP address changes occur.

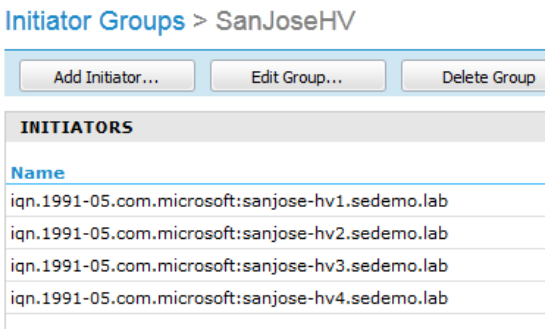

You should also select the "Allow multiple initiator access" when provisioning the Nimble volume which is a prerequisite for using SAN storage with a Microsoft Cluster. This option allows all nodes of the cluster to see the same volumes which permits high availability fail-over if a node fails.

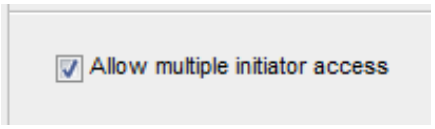

Connect the volume to each cluster host using the Microsoft iSCSI initiator. If you have multiple NIC ports from your host to the Nimble Storage array then you can enable MPIO to take advantage of the additional bandwidth and to make you storage connectivity more resilient to a single path failure. Once the volume is connected to each node, then using a single cluster node you can use the Disk Administrator on Windows 2008 R2 or File and Storage Services on Windows Server 2012 to see your new disk (Nimble volume) and bring it online by right-clicking on the disk. The figure to the right shows where to click using Disk Administrator in Windows 2008 R2.

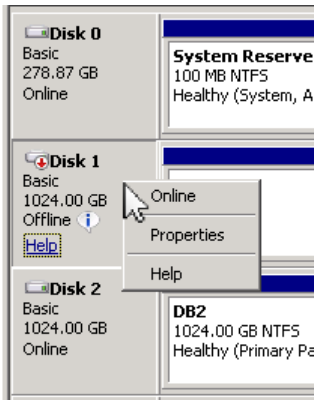

Windows Server 2012 doesn't show the Disk

If the disk that you have just attached using the Microsoft iSCSI Initiator isn't displayed immediately after connecting then you may need to refresh the Server Manager Disks view. Click the Refresh button located in the upper right portion of the Server Manager GUI and it will begin to re-inventory the server. The process completes when the scrolling status indicator in the title bar and the top of the Disks pane stops moving. You may then need to scroll the Disk list to find the new disk.

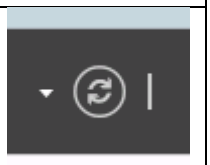

Determining Windows Disk Numbers

If you have many volumes attached to your system and you are not certain which disk is the volume that you're looking for then you can verify the Disk number using the Microsoft iSCSI Initiator. Select the volume in the targets list and click on the Devices button. The "Name" column will tell you the Disk number.

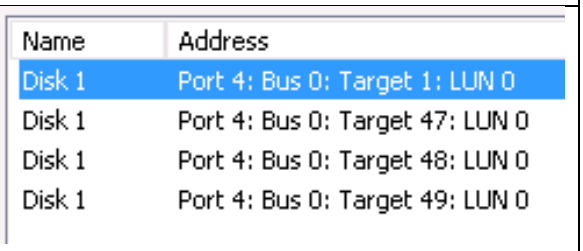

Name	Address
Disk 1	Port 4: Bus 0: Target 1: LUN 0
Disk 1	Port 4: Bus 0: Target 47: LUN 0
Disk 1	Port 4: Bus 0: Target 48: LUN 0
Disk 1	Port 4: Bus 0: Target 49: LUN 0

Disk Numbers are Not Always the Same

Windows Disk Numbers are not permanently assigned and may change between server reboots. In-addition the Disk Number may not match on each cluster node that mounts the same volume. This is by Windows design and does not affect cluster functionality.

 Use GPT Partitioned Windows Disks
When initializing a new disk you should prefer to initialize the partition table as GPT which is easier to work with than MBR volumes in many circumstances.

 Name Windows Volumes and Clustered Disks to match the Nimble Volume name
When possible, try to name your volumes and clustered disks the same as the underlying Nimble volume name. This best practice will prove useful when managing systems with many different disks and when working with clustered disks and cluster shared volumes.

The next step is to format the new disk as an NTFS volume. You do not need to specify a Drive Letter or Folder to mount the disk at this time. Use the default cluster size or force it to 4 KB for the CSV volume since it will hold operating system virtual disks and not data.

Once the Nimble volume is attached as a disk to each cluster node and formatted, you need to put it under control as a cluster resource to permit monitoring and fail-over between nodes. Start the Failover Cluster Manager tool in Windows, expand the Cluster and Storage trees.

Windows 2008 R2: Right click the Storage item and select Add Disk.

Windows Server 2012: Right-click on the Disks sub-item and select Add Disk.

Once the disk is added it will be assigned a name like “Cluster Disk #”. You should rename this clustered disk resource to match the Nimble volume name to help when managing your Hyper-V cluster. To rename the clustered disk, right-click on it and select properties, then you can edit the Name field.

Continue through the next section to configure the clustered disk as a Cluster Shared Volume.

Use Cluster Shared Volumes (CSV)

Windows 2008 R2 and 2012 Failover Clustering provides a new feature called Cluster Shared Volume (CSV) that provides an abstraction layer between the clustered application and the storage. CSV allows all Hyper-V nodes of the cluster to see the storage simultaneously, reducing the amount of time required for application failover and permitting storage of more than one virtual machine per iSCSI volume even if they are running on different nodes. You should enable cluster shared storage on your cluster by selecting the cluster item in the Failover Cluster Manager and then clicking the link in the Configure section (see screen shot).

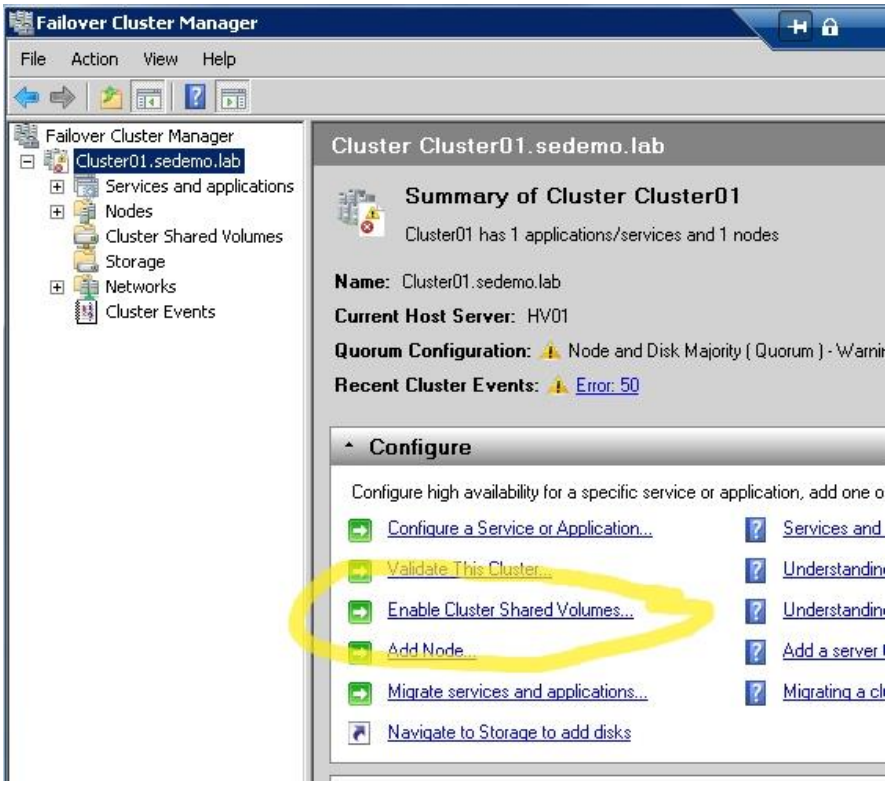

Windows Server 2008 R2

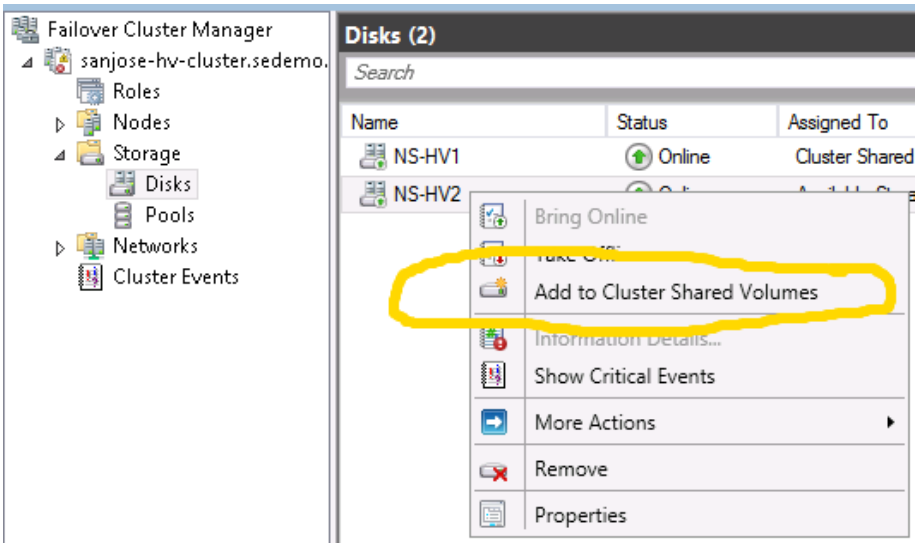

Windows Server 2012

Once Cluster Shared Volumes are enabled then you will see a new container in Failover Cluster Manager called “Cluster Shared Volumes”. CSV is implemented by mounting storage to each cluster node as junction points beneath the C:\ClusteredStorage directory. CSV creates a new sub-directory based on the format *Volume#* where # is a number that is incremented for each

successive volume attached as a CSV disk. If you have multiple Clustered Shared Volumes then you should take care to ensure that you're using the correct mount point and thus the correct CSV. If you're unsure about which mount point for a particular CSV, then you can verify it in Failover Cluster Manager:

Windows 2008 R2: Expand the CSV entry tree and the branch will display the mount point.

Windows Server 2012: Select the CSV Disk and the details pane below will display the mount point.

Use Protection Templates

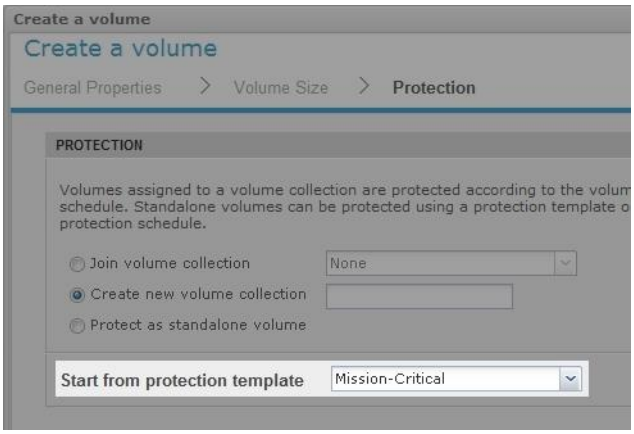

Nimble Storage arrays provide Protection Templates that consist of pre-configured schedules for snapshots, replication, and retention policies. When creating a new volume collection you can select a protection template that will insert a default schedule based on existing business rules. For example, you could create protection templates based on the criticality of the application data. Less critical applications can use longer snapshot schedule intervals (4 hours) and shorter retention schedules (10 days). However, more critical applications whose data frequently changes such as databases will usually require shorter snapshot schedule intervals (15 minutes or less) and longer retention schedules (90 days). Thus you will want to use a different protection template with shorter snapshot schedules and longer retention schedules. Using Protection Templates will reduce the amount of work required to create storage volumes and provide consistency for managing similar applications.

Use Hardware Snapshots versus Software Snapshots

Snapshots are the basis for creating point-in-time versions of storage volumes and backups that can be mounted and accessed just like any other iSCSI volume. You can create snapshots at different layers of virtualization architectures including within the Guest Software, within the Host Software, and within the Storage Hardware. Connecting data volumes directly to the guest allows NPM to trigger snapshots that use the Nimble hardware provider rather than inefficient software-based snapshots.

Nimble Storage arrays provide highly efficient hardware snapshot functionality that is optimized by Nimble's inline compression and block incremental efficiencies. This differs from operating system native software snapshots such as Microsoft™ VSS, which are not efficiently stored within their volumes. Thus software snapshots don't take advantage of Nimble Storage array optimized

snapshot backup functionality. The following diagram shows the differing locations in which snapshots are stored. It is preferable to use hardware-based snapshots in the Nimble Storage array that take advantage of performance, in-line compression, and cloning capabilities rather than performing software snapshots with far less flexibility.

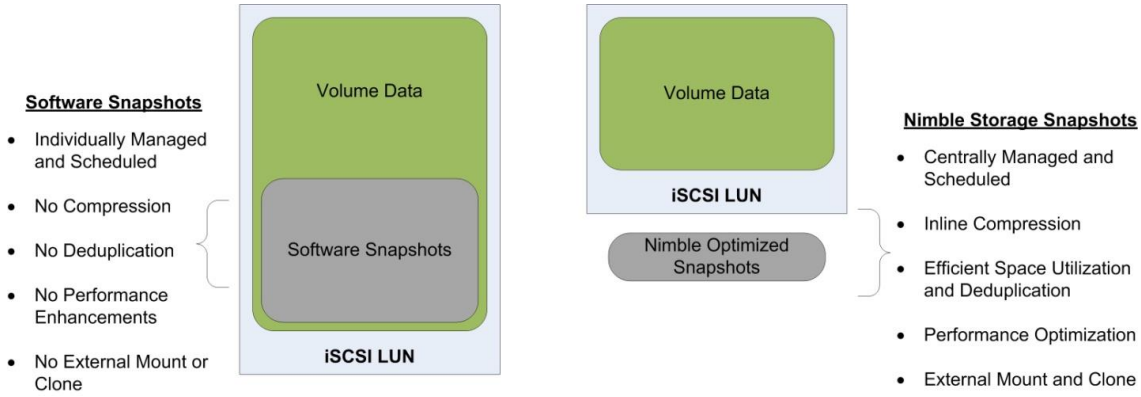

Use Zero-Copy Clones

Nimble Storage provides a feature called Zero-Copy Cloning that provides the ability to quickly clone a volume without duplicating all of the blocks of that volume. Zero-copy clones are valuable for creating test copies of production data without doubling the amount of space required to copy the data. You can clone production data volumes and mount them to test machines to perform Q/A testing and development, thus giving the benefit of working with production data and avoiding the chance of corruption. Another great use case for Zero-Copy Cloning is for reporting servers, this allows you to shift the ad-hoc reporting load off of production database servers without increasing the amount of storage space to hold it.

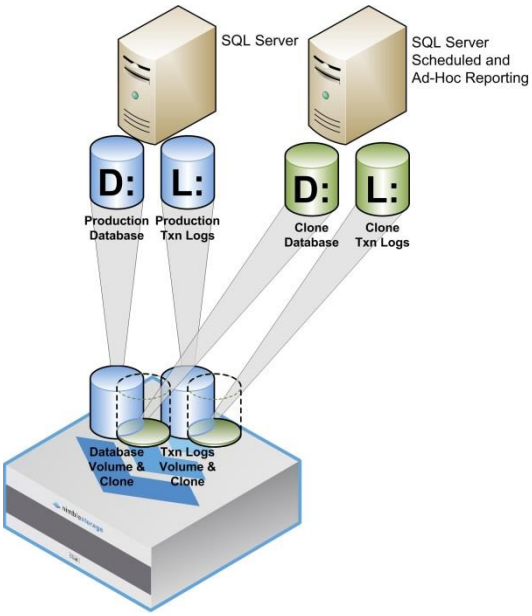

Provisioning Virtual Machines

There are three primary methods of provisioning virtual machines, using the native Hyper-V Manager , using the Failover Cluster Manager and using System Center Virtual Machine Manager. Hyper-V Manager can be started directly or used indirectly using the Failover Cluster Manager. You should avoid using the Hyper-V Manager directly since it will not create virtual machines on an individual server and not as a highly-available clustered resource. The Failover Cluster Manager will create a highly-available clustered virtual machine that can failover between cluster nodes.

Provisioning VMs using Hyper-V Manager and Failover Cluster Manager

Using either the Hyper-V Manager or the Failover Cluter Manager requires knowledge of where the CSV volumes are mounted to the file system. If you are unsure, then review the section *Using Clustered Shared Volumes (CSV)*. Once the virtual machine is created on the CSV then it will be cluster aware and permit failover between cluster nodes.

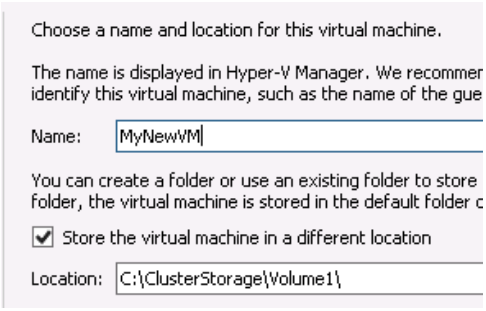

Provisioning VMs using System Center Virtual Machine Manager (SCVMM)

When provisioning virtual machines on your Hyper-V cluster with Nimble Storage using SCVMM it is important to enable High Availability or the virtual machine will not failover within the cluster properly. This feature is located in the Advanced section of the Configure Hardware step in the Create Virtual Machine Wizard. You can scroll down to find the Advanced section, once selected you will see the checkbox labeled “Make this virtual machine highly available”. By checking this box you will be able to place the virtual machine on the Nimble shared storage in the Configure Settings step of the Create Virtual Machine Wizard. Failing to select this option will only allow local non-shared storage for provisioning virtual machines.

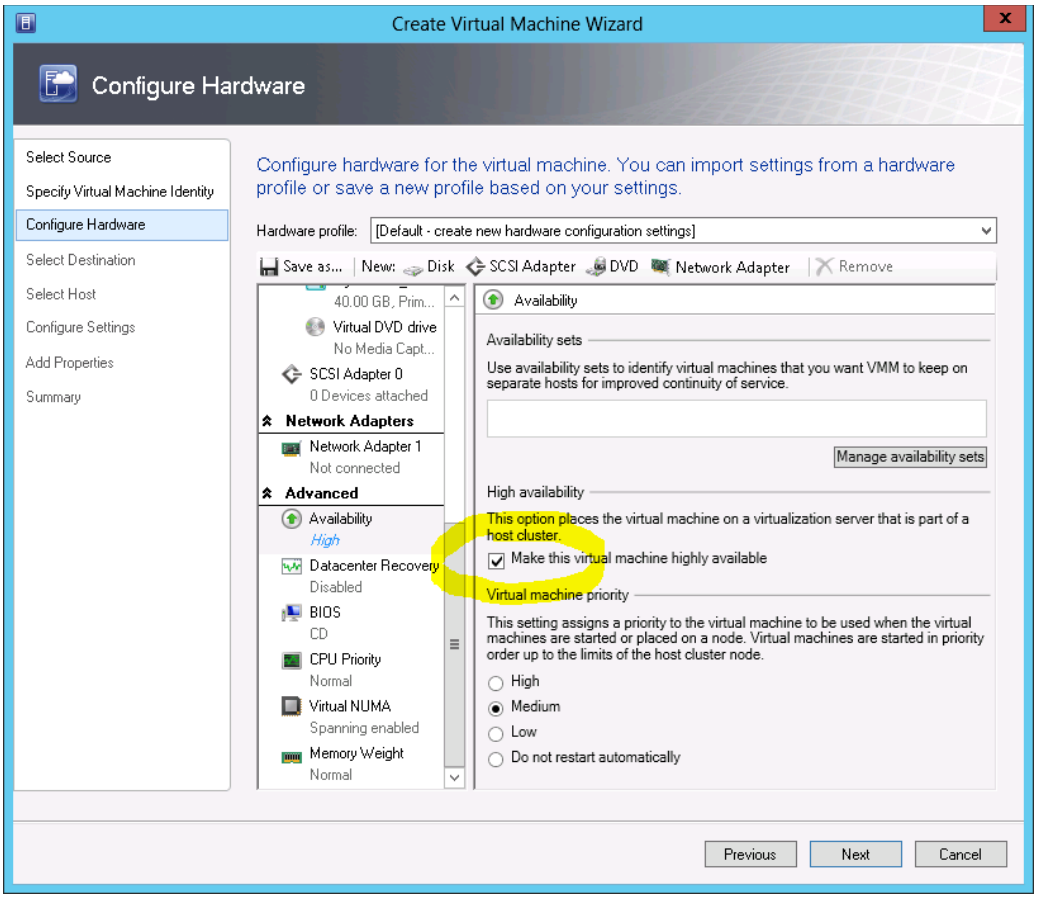

Select High Availability Option

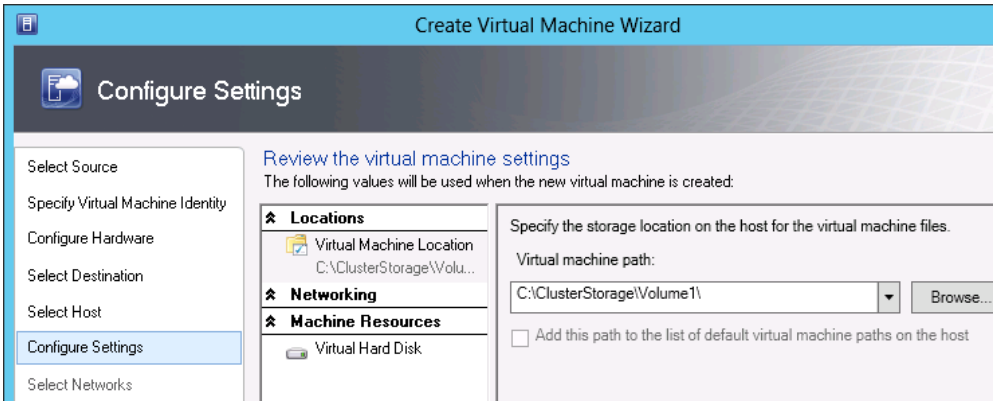

Specify CSV Storage

Better Hyper-V Backups

Protecting servers and data are primary goals for all IT administrators. Traditional methods required installing backup agents on each machine and then scanning the file system or application data to find data that has changed. Data size continues to grow and continues to put strain on the network and decrease backup windows. Both tax the production system resources during the backup process. In addition, the ease of virtually provisioning new servers has created the new phenomena of virtual server sprawl which adds to the growing problem of how to efficiently backup your servers and data. Providing better backup was a core founding challenge that Nimble Storage was created to solve. Nimble combines primary and backup storage into the same architecture and so avoids taxing the network to perform backup to backup storage. Nimble's highly efficient snapshot backup also allows you perform full backup much more frequently than using traditional backup technologies. This greatly improves your recovery point objectives and provides you with a true 24/7 backup window.

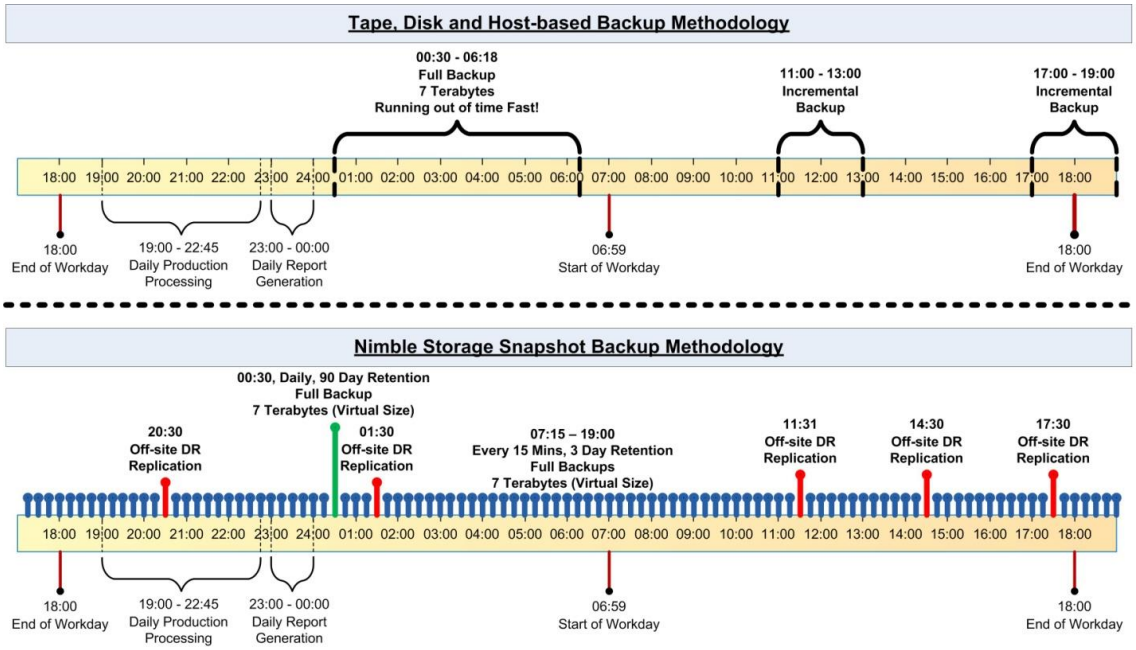

When a Nimble Volume Collection's protection schedule is triggered, the Nimble Protection Manager connects directly to the virtual machine's storage interface and asks it to place the application's data into a quiescent state. Applications begin to quiesce by flushing any pending I/O write activity from memory to disk and then signal NPM when they are ready for a safe snapshot backup. When NPM receives the quiesce notification, it triggers the Volume Collection to snapshot all its associated volumes, immediately after which data write activity is allowed to proceed. The Nimble backup method is dramatically faster and can trigger at regular short intervals unlike other solutions that have long backup windows that can take hours to complete before another backup can take place. Nimble Storage arrays perform snapshot backups instantly and can be scheduled for many more point-in-time backups per day than tape, disk, and Hyper-V host-based backup solutions. This is a big improvement over traditional backup, which leads many administrators to find that their backup windows continue to grow until they can no longer complete a daily backup, even with a 12-14 hour backup window. In addition, scheduled incremental backups leave gaps in protection and don't provide replication for off-site disaster recovery.

Off-Site Hyper-V Replication and Disaster Recovery

Nimble Storage arrays were built to replicate application-aware snapshots to other arrays and even off-site using a WAN-efficient methodology. Replication is configured on an individual virtual machine basis which allows you to choose which VMs that you need to protect based on their unique service level requirements. This also works well within the Nimble Storage Best Practices for Hyper-V framework which recommends using one VHD per Volume to take advantage of

Nimble's Zero-copy cloning features. There are two failover scenarios to consider when performing disaster recovery.

- **Planned Failover** – This disaster recovery scenario occurs when an outage is planned such as site maintenance or predictable such as a hurricane. Failover for this type of scenario is typically graceful by allowing you to shutdown applications and perform a final push of data from the production site to the disaster recovery site prior to starting the applications off-site.
- **Unplanned Failover** – This scenario occurs without prior notice and usually involves a site outage rather than a server failure. These are scenarios when you don't have time to perform a graceful failover such as the case of a fire or other catastrophic event. Failover may involve using older versions of volumes and losing data depending on your replication schedules. Thus you should plan your replication intervals shorter for more critical applications such as databases to reduce the potential data loss during unplanned failover.

Hyper-V DR Architecture

The disaster recovery architecture should use the same best practices as the production environment (specifically as Failover Clustering). However, the number of hosts does not need to match production precisely depending on your service level agreements. You should review the following considerations to ensure that your disaster recovery environment is able to perform failover properly using this implementation framework before referring to Appendix C to perform the disaster recovery steps.

Hyper-V Configuration Considerations

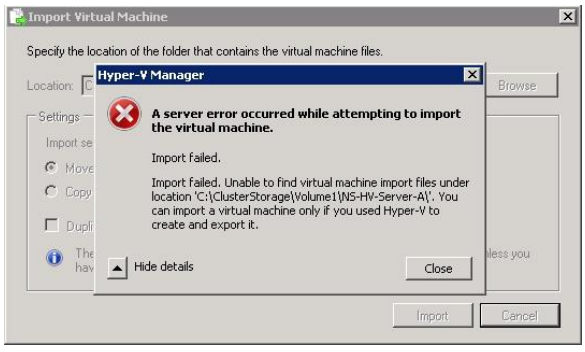

Microsoft Hyper-V R2 currently does not permit you to import virtual machines that were not previously exported by Hyper-V. Importing Hyper-V virtual machines can still be accomplished using a script to create the appropriate associations within the DR Hyper-V server. The script should be run for each virtual machine that you want to bring on-line in the disaster recovery site. Refer to Appendix C to create the Import-VM script that is used during Disaster Recovery failover and failback.

Virtual Network Naming Considerations

You should name your disaster recovery virtual networks with the exact same names as the production site to provide easier management. Hyper-V refers to virtual networks and attaches them to NICs using a unique ID which will be different in the disaster recovery site. Therefore, keeping consistent naming will assist you to reconnect virtual networks to the appropriate NIC. For example, if the public-facing NICs of your production virtual machines are attached to a virtual network name “Public vLAN”, then you should name the corresponding DR virtual network “Public vLAN”.

Restoration and Planned Failback

Once you have successfully failed over to your disaster recovery site and resumed business operations, new data will be created and modified over time. Restoring the new data back to your production site follows the same process as planned failover disaster recovery process in reverse. This will synchronize data back to the production array and allow you to resume business processes in your production facility. If you were forced to perform an unplanned failover, then you will need to begin the resynchronization manually by logging into your production array and selecting the volume collections that you failed over and clicking the Demote button. Then you must re-enable replication for each of the volume collections that you failed over on the disaster recovery array, which you can do while the volumes remain live.

Appendix A: Recovering a VM (Windows Server 2008 Hyper-V R2) using Import-VM Script

Background

This script is used to perform Hyper-V R2 failover since it does not have a native import feature that works without prior export of the virtual machine. This presents a Hyper-V management challenge to most Hyper-V environments, but can be overcome by registering the recovered virtual machine volumes and configuration into Hyper-V. The following code should be copied into a batch file called Import-VM.bat, which you store in both your production and disaster recovery environments.

```
@echo off

mklink "%systemdrive%\programdata\Microsoft\Windows\Hyper-V\Virtual Machines\%1.xml"
"%2\Virtual Machines\%1.xml"

icacls "%systemdrive%\programdata\Microsoft\Windows\Hyper-V\Virtual Machines\%1.xml" /grant
"NT Virtual Machine\%1":(F) /L

icacls %2 /T /grant "NT Virtual Machine\%1":(F)

rem done
```

Example Usage

You will need to determine the GUID for the Hyper-V machine that you want to import. This is located in the Virtual Machines directory in the VM volume's hierarchy and is the name of the XML configuration file. Ex. C:\ClusteredStorage\Volume1\<VM Name>\Virtual Machines\ 58AE37DA-53A1-412F-996E-9E26C602696D.xml"

Import-VM <GUID> "<Path to Config>"

Import-VM 58AE37DA-53A1-412F-996E-9E26C602696D "C:\ClusteredStorage\Volume1\NS-HV-Server-A"

Note the quotes surrounding the path to the configuration file. After running the script, you will see output such as the following. Note that the first line run creates a symbolic link which reports success as "created". Next permissions are changed for the directories to permit the virtual machine's identity to connect to its' configuration files, the last line reports that no files failed processing.

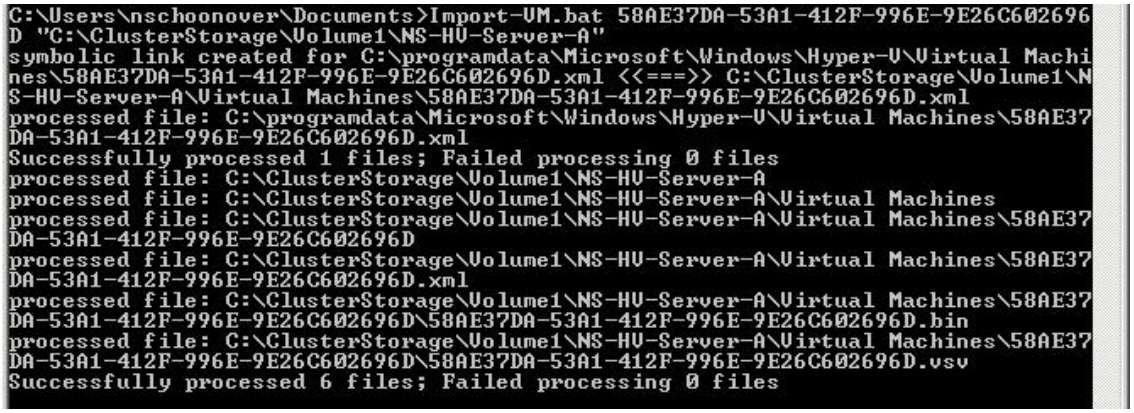

If the script fails initially, then verify the input parameters and run again. If the script continues to fail and you want to run from a clean point, then you can begin fresh by deleting the symbolic link for the VM GUID in the normally hidden directory “C:\ProgramData\Microsoft\Windows\Hyper-V\Virtual Machines”.

Appendix B – Disaster Recovery Failover (Windows Server Hyper-V R2)

Use the following steps to perform either planned or unplanned disaster recovery failover using Nimble Storage replicated volumes.

Planned Failover

1. Handover production volumes.

   1.1. Gracefully shutdown the applications and virtual servers that you want to failover.

   1.2. Login to your production Nimble Storage array.

   1.3. Select Manage -> Protection from the menu to view the Volume Collection.

   1.4. Select the Volume Collection that you want to failover.

   1.5. Click the Handover button. The Handover process will take a snapshot of the volumes and begin copying the most recent data changes to the disaster recovery array. This may take some time depending on how much data has changed.

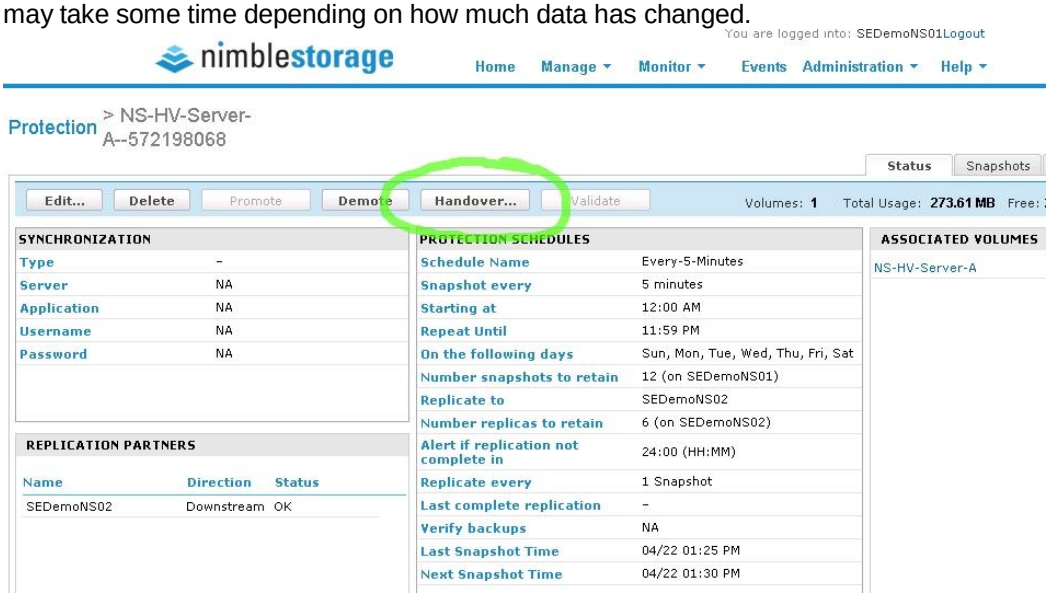

   1.6. Handover is complete when the Volume Collection icons and status on the disaster recovery array change as in the following graphic. You may have to refresh the Volume Collection view by selecting the Home page and then selecting the Manage -> Protection page in the Nimble user interface.

Unplanned Failover

1. Promote volumes

    1.1. Login to your disaster recovery Nimble Storage array.

    1.2. Select Manage -> Protection from the menu to view the Volume Collection.

    1.3. Select the Volume Collection that you want to failover.

    1.4. Click the Promote button, which is used only for failover when the production array is no longer available such as in the case of an unplanned failover. This will change the ownership of the volume collection to the disaster recovery array.

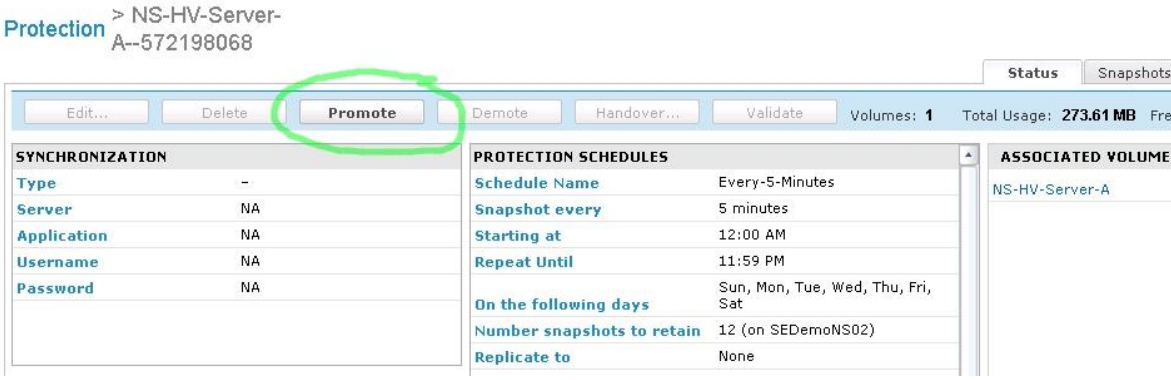

    1.5. The promotion is complete when the volume collection icon and status have changed similar to the Handover process described in Planned Failover.

Steps Shared by Failover Methods

2. Mount the disaster recovery volume

    2.1. Mount the volume to a Hyper-V host in the disaster recovery site using the Microsoft iSCSI Initiator.

    2.2. Bring the volume online using the Disk Administrator.

    2.3. Use Failover Cluster Manager to add the volume to the Storage container. To do this, right click on the Storage container and select Add, then select the volume in the dialog box.

    2.4. Next, add the disk to the Clustered Shared Volumes container by right clicking the container and selecting Add, then select the volume in the dialog box. This step will mount the volume to a junction point in the C:\ClusteredStorage directory.

3. Import the Hyper-V virtual machine

Windows 2012 Hyper-V (R3) Procedure

3.1. Use the Hyper-V Manager to right click on your server and select Import Virtual Machine.

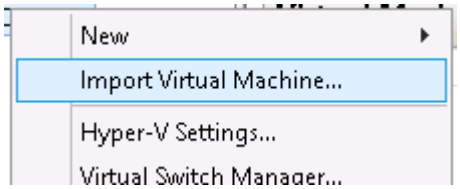

3.2. Select the folder containing the virtual machine in the Clustered Storage volume that you connected in the previous step.

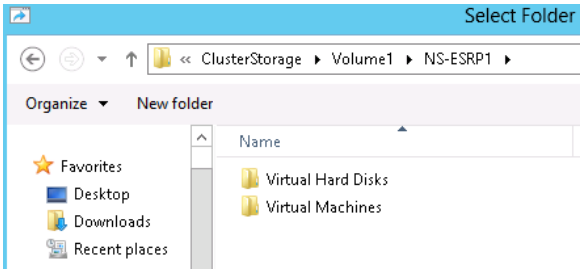

3.3. Follow the rest of the steps in the Import wizard to Finish importing the VM.

3.4. Continue adding the imported virtual machine as a clustered resource in step 4.

Windows 2008 R2 Hyper-V Procedure

3.1. Open Windows Explorer and go to the Virtual Machines directory located in the "C:\ClusteredStorage\<Imported Volume>\<Virtual Machine>\" sub-directory for the Nimble volume that you just added to Failover Clustering. You are looking for the Hyper-V configuration file that will have a name that represents the GUID of the virtual machine. A new configuration file and associated sub-directory are created every time that you create a virtual machine, thus if you see multiple configuration files and you are following Nimble best practices of one virtual machine VHD per volume then you should select the most recent configuration and can safely delete the older configurations. If you are hosting multiple virtual machines from the same volume, then you will need to run the Import-VM script for each virtual machine.

3.2.

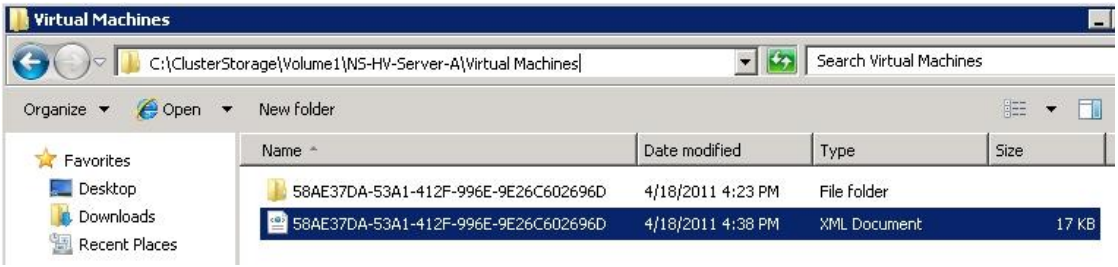

3.3. Select the GUID for the virtual machine to import. The easiest method is to select the file or directory and then clicking them to rename the file and then right-clicking the highlighted text and select copy.

3.4. Next open a command line on the Hyper-V host that can run the Import-VM script that you should have created in Appendix C.

3.5. Run the Import-VM script, substituting the GUID and full path to the configuration file. After successful completion you may need to restart the Hyper-V service to see the virtual machine in the Hyper-V Manager. Using our example in the previous graphic the command would look like this:

Import-VM 58AE37DA-53A1-412F-996E-9E26C602696D
"C:\ClusteredStorage\Volume1\NS-HV-Server-A"

3.5. Continue adding the imported virtual machine as a clustered resource in step 4.

4. Add Imported Virtual Machine as Clustered Resource

4.1. After the virtual machine is properly imported into Hyper-V, then you should add it to the disaster recovery clustered applications. Using Failover Cluster Manager, right click on Services and Applications then select Configure a Service or Application. Note: DO NOT try to add the imported VM using the Virtual Machines sub-menu as those links will only allow the creation of a new virtual machine.

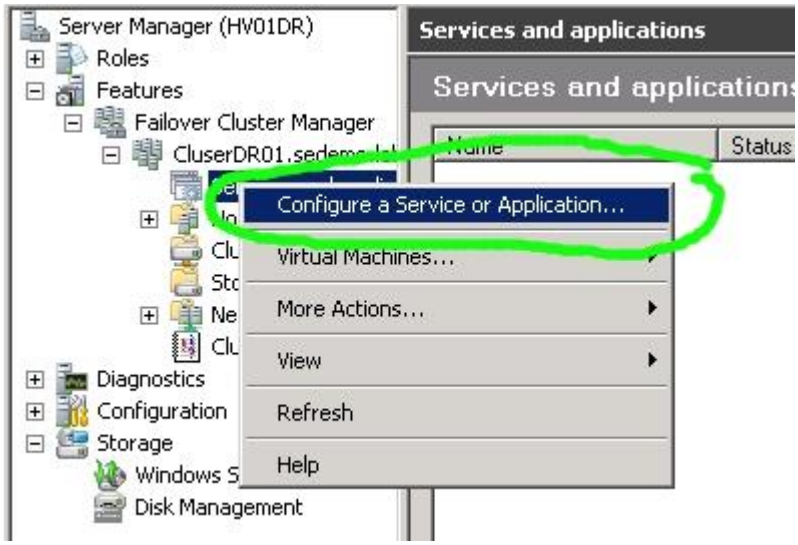

4.2. Select Virtual Machine as the resource type.

4.3. The next screen in the wizard should display the virtual machine that you have just imported. Select that checkbox and click the Next button.

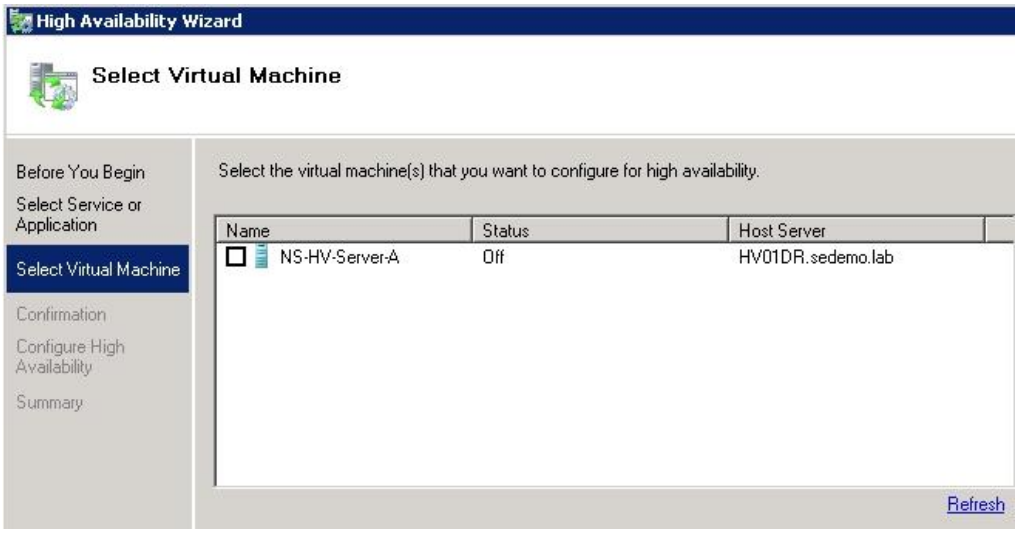

4.4. Finally, right-click on the VM in the Failover Cluster Manager and click Start virtual machines.

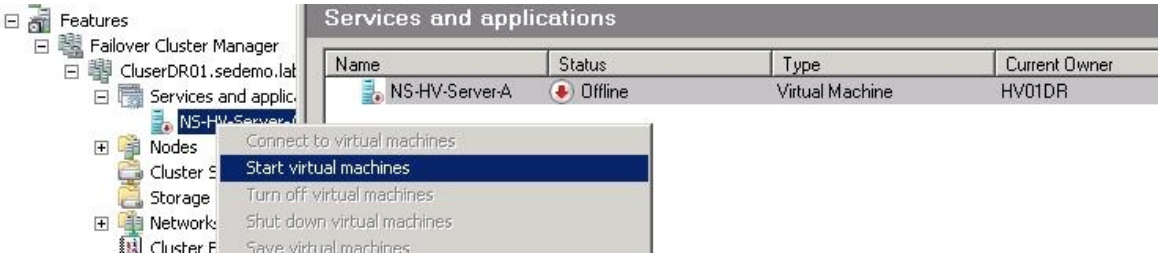

Nimble Storage, Inc.

2740 Zanker Road., San Jose, CA 95134

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