

BEST PRACTICES GUIDE:

Nimble Storage Best Practices for Microsoft Hyper-V R2



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 architecture using Nimble Storage and Windows 2008 R2 Hyper-V.

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.
- **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.
- **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.
- **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.
- **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.

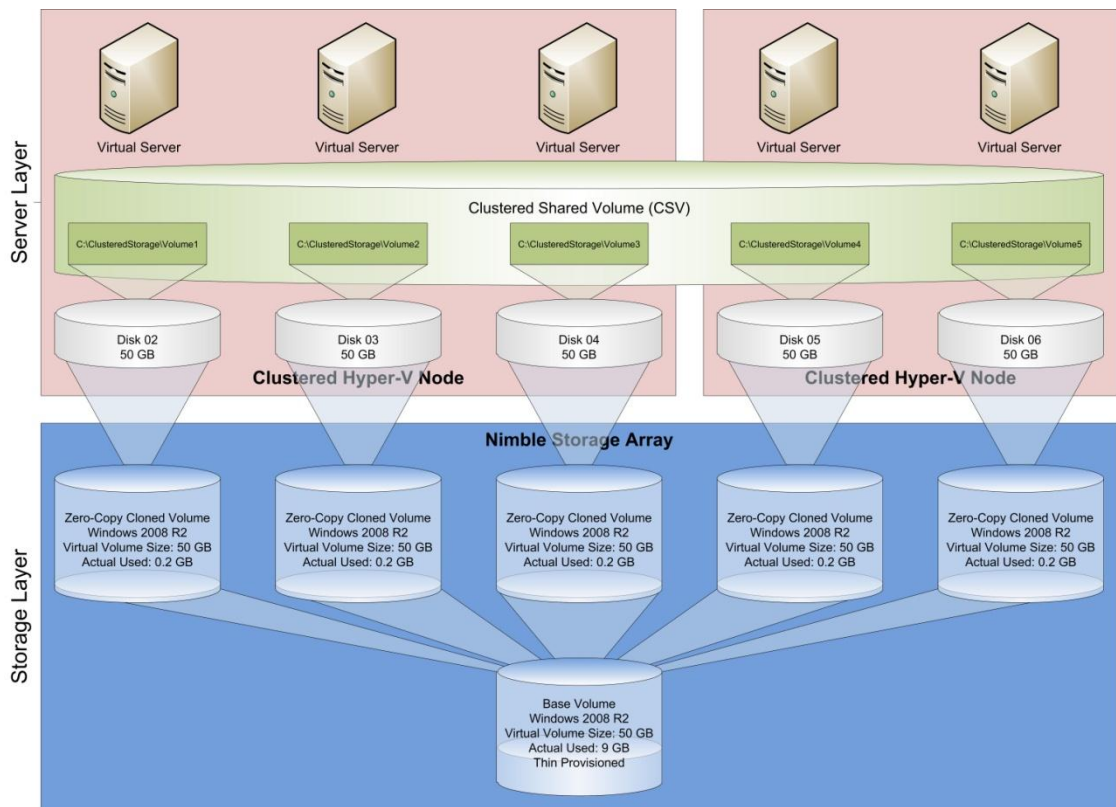
Hyper-V Storage Architecture

Implementing 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 your Hyper-V cluster nodes.

Nimble Storage Software
Required

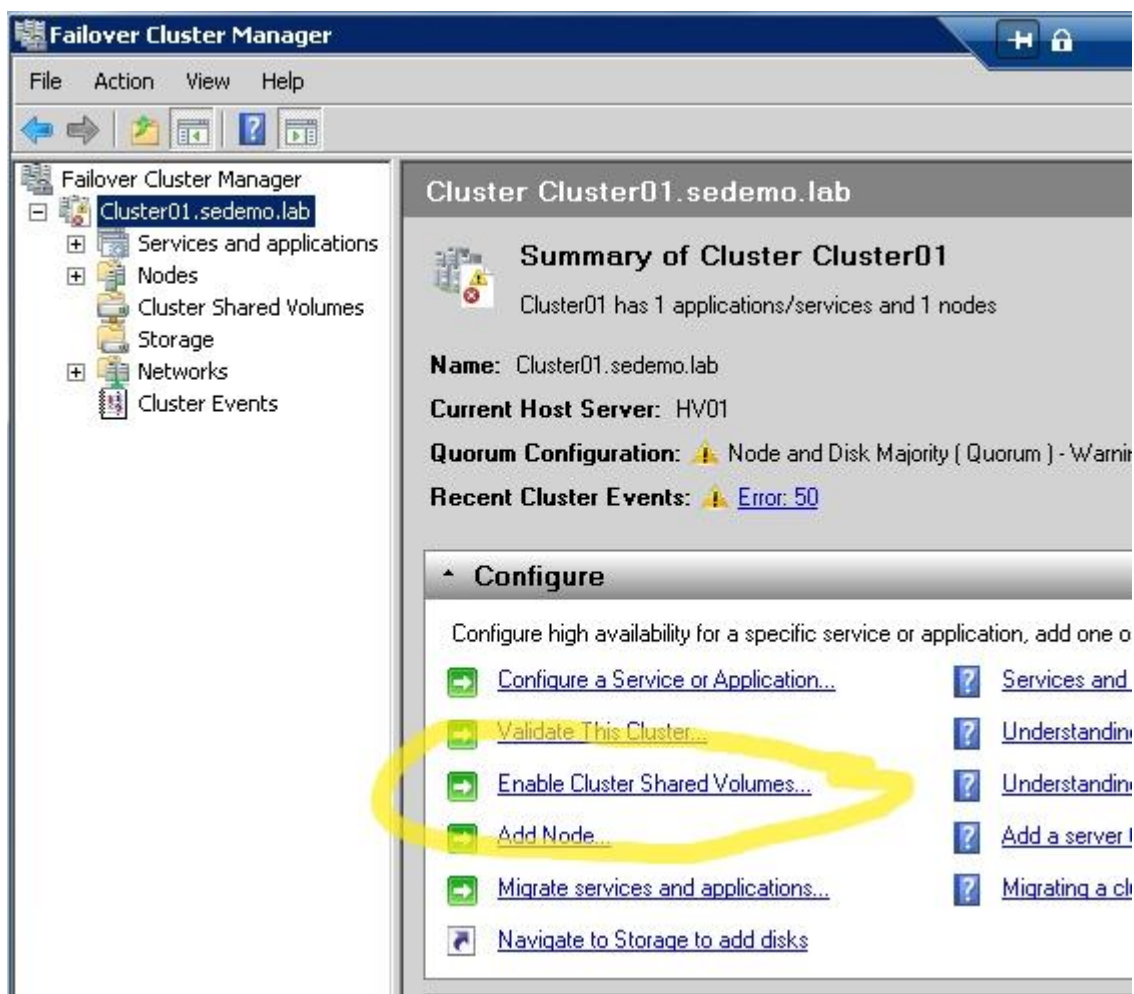
1.0.6.2 or higher

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



Use Cluster Shared Volumes

Windows 2008 R2 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. 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).



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.

Provision One Volume per Operating System Boot Image and CSV

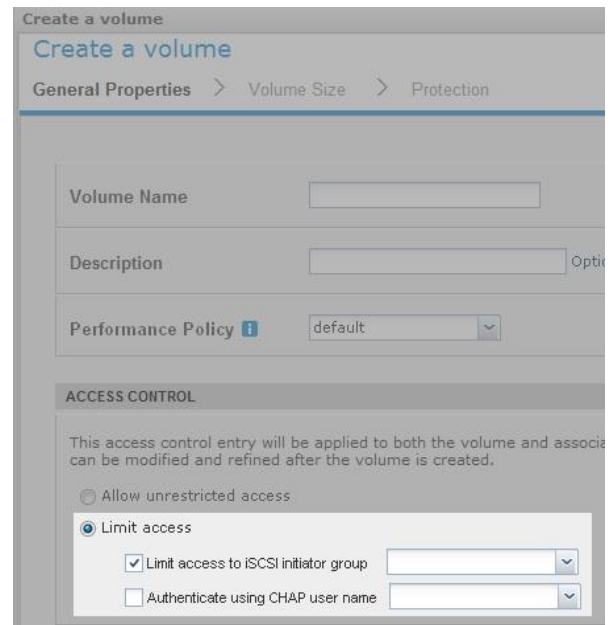
Provisioning storage for Operating System boot volumes can be done in a many-to-one or one-to-one configuration. Using a one-to-one ratio provides the maximum space savings when using Nimble Storage zero-copy clones and enables granular disaster recovery of applications.

Microsoft Failover Clusters use SCSI-3 Persistent Reservations to ensure that only the owning node can write to a disk. Many storage arrays have limits on the number of persistent reservations, which limit their effectiveness to support a clustered Hyper-V architecture. This forces you to use a many-to-one VHD-to-CSV configuration. Nimble Storage arrays have no SCSI-3 persistent reservation limit, which allows this clustered Hyper-V architecture to scale to hundreds of virtual machines.

Refer to the steps detailed in Appendix A to provision a new Nimble Storage volume for use as a Base Operating System boot image. We will also mount that volume to a Hyper-V host to install the OS on a Virtual Hard Disk (VHD) and SysPrep it for use with rapid virtual machine provisioning.

Limit Access to All Volumes

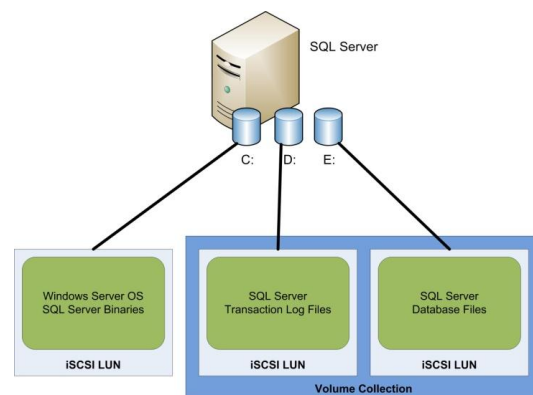
When using volumes in environments that run multiple hosts, it is important to isolate all volumes only to the hosts or virtual machines that should have access to them. Failing to limit access to iSCSI volumes will clutter user interfaces and cause some server platforms to perform excessive connections to all volumes—including those that it should not have access to—possibly placing your data at risk that another virtual machine could attach to them and delete data. Limiting access by Initiator Group greatly reduces this risk that can affect the performance of storage operations such as taking snapshots. Nimble Storage provides the ability to limit access to a specific iSCSI initiator group (a list of one or more iSCSI initiator IQN identifiers). You can also control access by using CHAP authentication to further secure your iSCSI volumes.



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.

It is also helpful to separate 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. 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



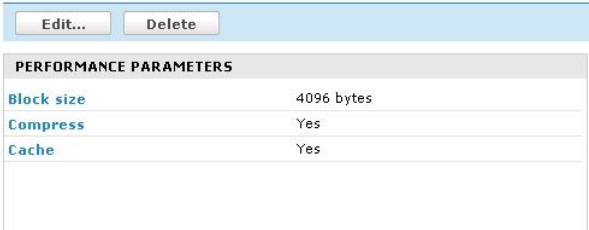
non-trusted, state that forces a complete restoration from a backup. Therefore, it is important to protect both the transaction logs and database in a coordinated fashion when performing any type of backup operation. Nimble Storage arrays provide functionality that allows you to group volumes that need to be mutually consistent into the same Volume Collection.

Use Performance Policies

The Nimble Storage array includes profiles called Performance Policies that pre-configure new volumes using optimized configuration settings specific for different usage scenarios. For example, the default performance policy is tuned to use 4 KB volume blocks to provide the best

performance for Windows storage LUNs. As you can see in the screen shot, the default performance policy also includes in-line compression and high-performance caching. Thus, use the default performance policy for Windows volumes. You should 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.

Performance Policies > default

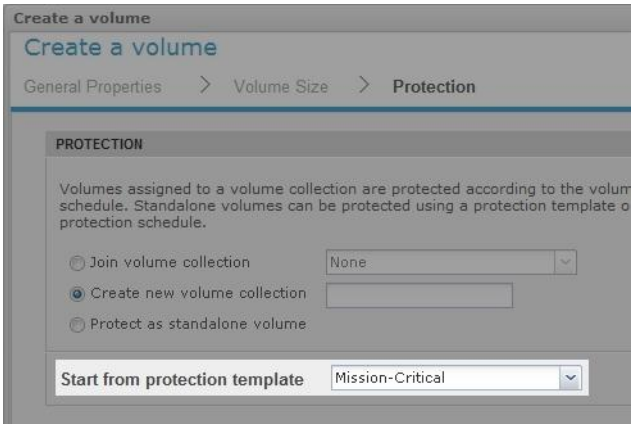


PERFORMANCE PARAMETERS	
Block size	4096 bytes
Compress	Yes
Cache	Yes

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.



PROTECTION

Volumes assigned to a volume collection are protected according to the volume schedule. Standalone volumes can be protected using a protection template or protection schedule.

Join volume collection None

Create new volume collection

Protect as standalone volume

Start from protection template Mission-Critical

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. Management of a volume collection allows you to quickly change protection

schedules for all related volumes. For example, you have created a SQL Server database protection schedule for several databases on a SQL Server supporting an eCommerce application. It is common for databases to be partitioned into different data files with differing performance characteristics. The initial business requirements called for a configuration based on hourly snapshots for backup and replication off-site every 6 hours for disaster recovery. As business has increased, management has decided that the database has become more critical and thus needs more frequent backup and more frequent replication to reduce potential data loss. You can change the protection schedule for all of the associated volumes for the database at the same time using a Volume Collection, thus saving time and eliminating configuration errors that might inhibit recoverability.

Edit Volume Collection

Introduction Synchronization Schedules

Snapshots are point-in-time copies of volumes that allow you to recover data taken with no performance impact and are very space efficient, allowing you to backup copies on the array. Snapshot schedules specify how frequently snapshots will be retained for the volumes assigned to this volume collection.

Schedule Name: Hourly

Repeat Every: 1 hours

Starting at: 12:00 HH:MM AM

Repeat Until: 11:59 HH:MM PM

On the following days: Mon Tue Wed Thu

Number of snapshots to retain: 1

Replicate to: None

Verify backups: Yes

Add another Schedule

Prefer Guest Connected iSCSI Volumes for Storing Data

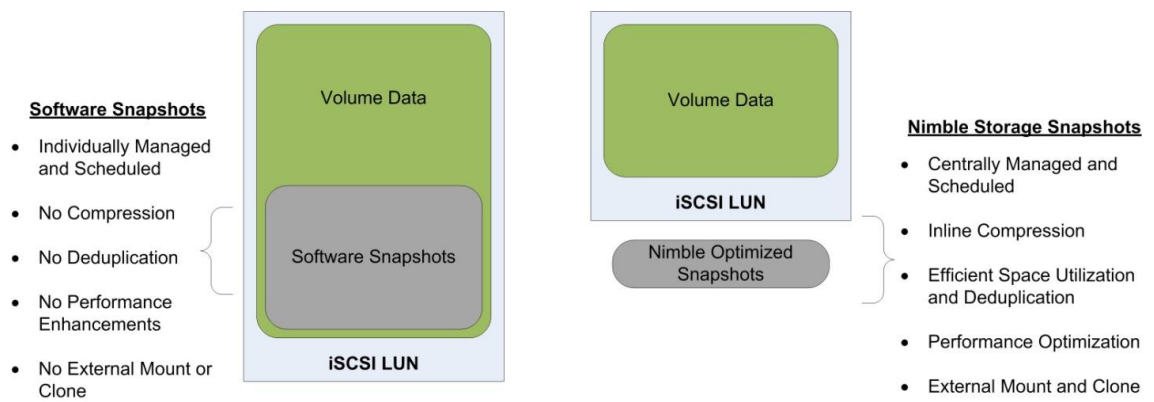
You should store data on Nimble Storage iSCSI volumes rather than through virtual hard disks. This method of data storage connectivity provides the best solution for data protection, application consistency, and off-site replication, as well as performance enhancements, by making complete use of Nimble Storage array optimization features. You should install Nimble Protection Manager (available for download from the Nimble Storage support web site) on each of your guest virtual machines that store data. You may even want to install NPM into the base image of your virtual machines and then enable it when needed.

Prefer 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 gracefully clone a volume without duplicating all of the blocks of that volume. This feature allows you to create a base OS image that is standard for all servers in your environment, complete with service pack, patches, and additional third-party software such as anti-virus; then you can create clones of that base image for each new virtual machine that you create. Using zero-copy cloning of Hyper-V OS volumes, as shown in the previous storage architecture diagram, saves 9+ GB of storage space for each virtual machine that you create by eliminating duplicated blocks that are common to all clones. Thus, we have effectively saved 36 GB of storage space in an environment hosting only five virtual machines. As the number of virtual machines grows, these savings begin to increase exponentially.

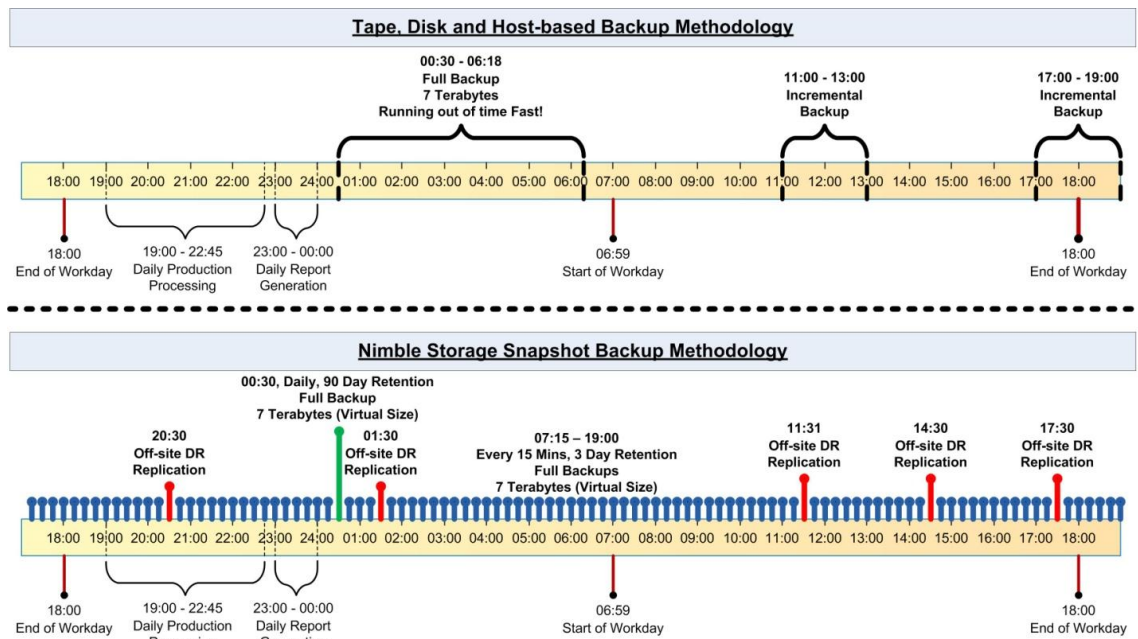
Virtual Machines	Storage Savings
5	36 GB
10	81 GB
25	216 GB
50	441 GB
100	891 GB

Zero-copy clones are also valuable for creating test copies of production data without making full copies of 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.

Refer to the steps detailed in Appendix B to create a Nimble Zero-copy clone of a base OS boot volume for use with rapid Hyper-V virtual machine provisioning.

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 VMware 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.

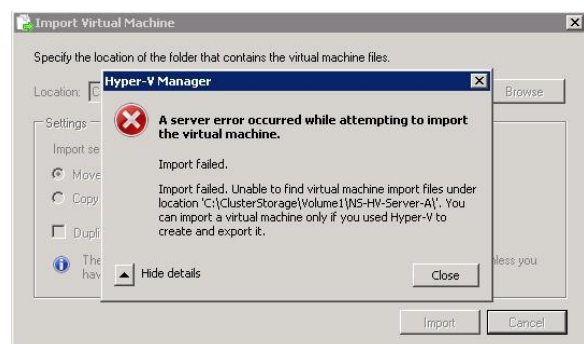
- **Planned Failover** – This disaster recovery scenario occurs when an outage is planned such as site maintenance or predictable such as a hurricane. Failover for this type of scenario is typically graceful by allowing you to shutdown applications and perform a final push of data from the production 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: Provisioning Base Volumes for Use in Cloning

1. Create a Volume
 - 1.1. Prefer a Volume name that will be obvious that it is an infrastructure volume, what it contains and won't be confused as a production data set. Ex. Base-HV-Win2k8R2, meaning that it's a Base infrastructure volume for Hyper-V (HV) and contains a Win2k8R2 SysPrepped image.
 - 1.2. Use the Default Performance Policy.
 - 1.3. Specify an iSCSI Initiator Group for your Hyper-V hosts.
 - 1.4. Enable "Allow multiple initiator access"
 - 1.5. Specify the volume size. Make sure to leave disk space for PageFile and the applications that you will be installing on the cloned volumes but not data. Recommend at least 50 GB. Note: Nimble will Thin Provision the volume by default, only consuming space as it is written.
2. Mount the Volume to the Hyper-V host
 - 2.1. Use the host's Microsoft iSCSI initiator and connect the volume. Click on the Advanced button and specify the path that the initiator should use to ensure that iSCSI traffic is forced to the data network.
 - 2.2. Attach additional connections specifying different initiator IP to Target IP mappings if using MPIO.
 - 2.3. Click on the attached volume and click the Devices button. Make note of the Disk number. Do not attach to other cluster nodes yet or the volume will be marked as a cluster volume, which we don't want at this point to create the initial SysPrepped image.
 - 2.4. Online the Disk.
 - 2.5. Initialize the Disk as GPT.
 - 2.6. Quick Format as NTFS. Do not use compression. Assign a drive letter or mount point.
3. Create a Virtual Machine in Hyper-V
 - 3.1. Use the Hyper-V Management Console, not the Failover Cluster Manager, to create the base virtual machine.
 - 3.2. Ensure that the VM Location is in the volume for the Nimble volume that was attached.
 - 3.3. Create a new Virtual Disk. Remember that Nimble uses Thin Provisioning by default. Fill the volume (49 GB) for the VHD.
 - 3.4. Attach an OS install ISO image to the CD-ROM/DVD drive.
 - 3.5. Start the virtual machine.
4. Install Operating System
 - 4.1. Install OS. Note: If you already have a base configuration Virtual Hard Disk then you can copy it to the Nimble volume and skip these steps, but you should start it to install the .NET 3.5 Feature and Nimble Protection Manager.
 - 4.2. Install Virtual Integration Tools if necessary.
 - 4.3. Install or enable the .NET 3.5 feature for Windows OS.
 - 4.4. Service Pack
 - 4.5. Apply Hot Fixes and Patches
 - 4.6. Install Nimble Protection Manager. This will allow machines that require application consistent snapshots by quiesce to properly coordinate between the application and Nimble Storage array.
 - 4.7. Install additional standard software that all virtual machines will have such as Anti-Virus.
 - 4.8. SysPrep the image using `C:\Windows\System32\SysPrep\Sysprep.exe /generalize /oobe /shutdown`
5. Finalize the cloned image

- 5.1. After the SysPrep completes and the VM is shutdown, then use the Hyper-V console and delete the virtual machine. Note: This will just delete the virtual machine configuration and will leave behind the SysPrepped virtual hard drive which we will use to clone and build future virtual machines.
- 5.2. Use the Hyper-V host's Disk Management tool to off-line the SysPrep base volume.
- 5.3. Use the Hyper-V host's iSCSI Initiator and disconnect the Nimble base config volume. You should also remove it from the Favorites tab.
- 5.4. Use the Nimble Storage GUI and take the SysPrepped volume off-line to ensure that the volume isn't altered.
- 5.5. Create a snapshot of the volume using the Nimble Storage GUI and label it as the SysPrepped image, also adding some text in the description specifying what the OS image contains. You should use a description that contains the base volume name. New clones created from this snapshot will copy the description which will allow you to know which base volume the clone was derived from.

Appendix B: Using Nimble Clones for Rapid Hyper-V Provisioning

1. Select a SysPrepped volume.
 - 1.1. Use the Nimble GUI and select the SysPrepped image that you wish to clone from the volumes list.
 - 1.2. Select the Snapshots tab and click the check box next to the snapshot for the version of the image that you wish to use, perhaps the only one.
 - 1.3. Click the Clone button and specify the name of the new cloned volume.
 - 1.4. Select the new cloned volume. Ensure that it has the iSCSI Initiator Group of the base volume. Click the Set Online button
2. Mount the Volume to the Hyper-V Host
 - 2.1. Use the host's Microsoft iSCSI initiator and connect the volume. Click on the Advanced button and specify the path that the initiator should use to ensure that iSCSI traffic is forced to the data network.
 - 2.2. Attach additional connections specifying different initiator IP to Target IP mappings if using MPIO.
 - 2.3. Click on the attached volume and click the Devices button. Make note of the Disk number. Do not attach to other cluster nodes yet or the volume will be marked as a cluster volume, which we don't want at this point to create the initial SysPrepped image.
 - 2.4. Online the Disk.
 - 2.5. Change the volume ID using DiskPart
 - 2.5.1. Open a command prompt with Administrator privileges and type DiskPart.
 - 2.5.2. Select the disk that you just mounted using the DiskPart command Select Disk #. If you don't know the Disk number then you can see it in the Windows Disk Manager or by typing List Disk.
 - 2.5.3. Remove the ReadOnly attributes. Type Attributes Disk Clear Readonly.
 - 2.5.4. Type UniqueID Disk to view the current disk ID which will allow Windows to mount multiple cloned volumes. Change the Unique ID to a value other than the existing Unique ID by typing UniqueID Disk id=#. Note: This is a GUID for GPT volumes and a Hexidecimal number for MBR volumes. You can use <http://www.newguid.net> to generate a new GUID.
 - 2.5.5. Type Exit to leave DiskPart.
3. Configure MSCS CSV Storage
 - 3.1. Use the MSCS management console, select the Storage container. Right-click and select "Add a Disk". The Nimble volume should be visible, select just that disk and press OK.
 - 3.2. Repeat the previous step for each member host server of the Hyper-V Cluster.
 - 3.3. Select the Cluster Shared Volumes container. Right-click and select "Add Storage", specifying the Storage volume that you just added. This will make the volume available to all clustered Hyper-V hosts simultaneously with proper coordination. Note: CSV volumes are mounted as junction points below the C:\ClusteredStorage directory and have a directory format as Volume#. The Volume# is created in sequence when the disks are originally attached and do not match the physical Disk # in Windows Disk Manager.
4. Create a Virtual Machine
 - 4.1. Select the Services and Applications container. Right-click and select Virtual Machines, New Virtual Machine and the host that you want to initially deploy the VM.
 - 4.2. Ensure that the VM Location is in the CSV directory for the Nimble volume that was attached. Note: Hyper-V GUI tends to display previously used paths by default, ensure that you are placing the VM in the correct junction point under C:\ClusteredStorage.

- 4.3. Select the previously created VHD in the appropriate C:\ClusteredStorage junction point rather than creating a new VHD.
5. Perform any additional VM configuration using the Hyper-V Management console. I.e. Additional NICs. Note: Do not add connections for data that will be stored on Nimble, as all data should be connected directly via Guest iSCSI Initiator to ensure proper consistency through quiesce.
6. Start the Virtual Machine

Appendix C: 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.

```
C:\Users\nschoonover\Documents>Import-UM.bat 58AE37DA-53A1-412F-996E-9E26C602696
D "C:\ClusterStorage\Volume1\NS-HU-Server-A"
symbolic link created for C:\programdata\Microsoft\Windows\Hyper-V\Virtual Machi
nes\58AE37DA-53A1-412F-996E-9E26C602696D.xml <<==>> C:\ClusterStorage\Volume1\N
S-HU-Server-A\Virtual Machines\58AE37DA-53A1-412F-996E-9E26C602696D.xml
processed file: C:\programdata\Microsoft\Windows\Hyper-V\Virtual Machines\58AE37
DA-53A1-412F-996E-9E26C602696D.xml
Successfully processed 1 files; Failed processing 0 files
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A\Virtual Machines
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A\Virtual Machines\58AE37
DA-53A1-412F-996E-9E26C602696D
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A\Virtual Machines\58AE37
DA-53A1-412F-996E-9E26C602696D.xml
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A\Virtual Machines\58AE37
DA-53A1-412F-996E-9E26C602696D\58AE37DA-53A1-412F-996E-9E26C602696D.bin
processed file: C:\ClusterStorage\Volume1\NS-HU-Server-A\Virtual Machines\58AE37
DA-53A1-412F-996E-9E26C602696D\58AE37DA-53A1-412F-996E-9E26C602696D.vsv
Successfully processed 6 files; Failed processing 0 files
```

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 D – Disaster Recovery Failover

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.


The screenshot shows the Nimble Storage web interface. At the top, the logo and navigation menu are visible. The main content area is titled 'Protection > NS-HV-Server-A--572198068'. Below the title, there are several tabs: 'Edit...', 'Delete', 'Promote', 'Demote', 'Handover...', and 'Validate'. The 'Handover...' button is highlighted with a green circle. Below the tabs, there are three main sections: 'SYNCHRONIZATION', 'PROTECTION SCHEDULES', and 'ASSOCIATED VOLUMES'. The 'SYNCHRONIZATION' section shows details for the volume collection, including Type, Server, Application, Username, and Password. The 'PROTECTION SCHEDULES' section shows details for the protection schedule, including Schedule Name, Snapshot every, Starting at, Repeat Until, On the following days, Number snapshots to retain, Replicate to, Number replicas to retain, Alert if replication not complete in, Replicate every, Last complete replication, Verify backups, Last Snapshot Time, and Next Snapshot Time. The 'ASSOCIATED VOLUMES' section shows the volume collection name 'NS-HV-Server-A'.

- 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.

Volume Collection Icon and Status During Normal Replication

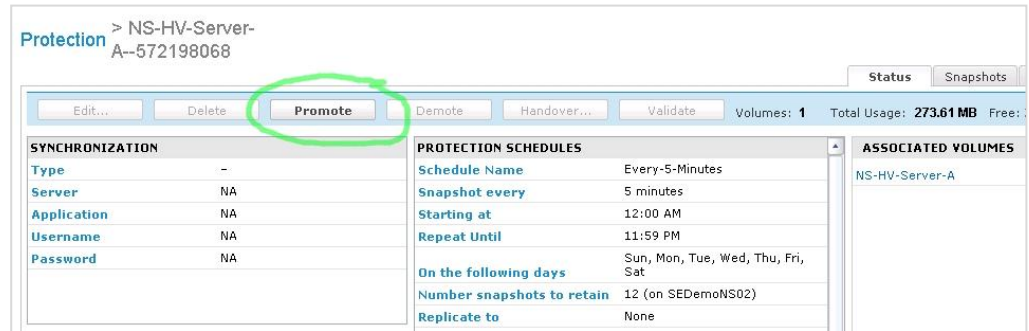
	NS-HV-Server-A--572198068	None	◀ SEDemoNS01
---	---	------	--------------

Volume Collection after Handover

	NS-HV-Server-A--572198068	None	▶ SEDemoNS01
---	---	------	--------------

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.

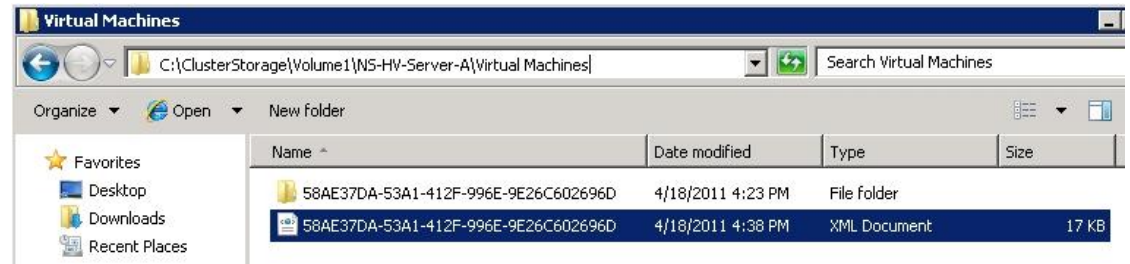


- 1.5. The promotion is complete when the volume collection icon and status have changed similar to the Handover process described in Planned Failover.

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
 - 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

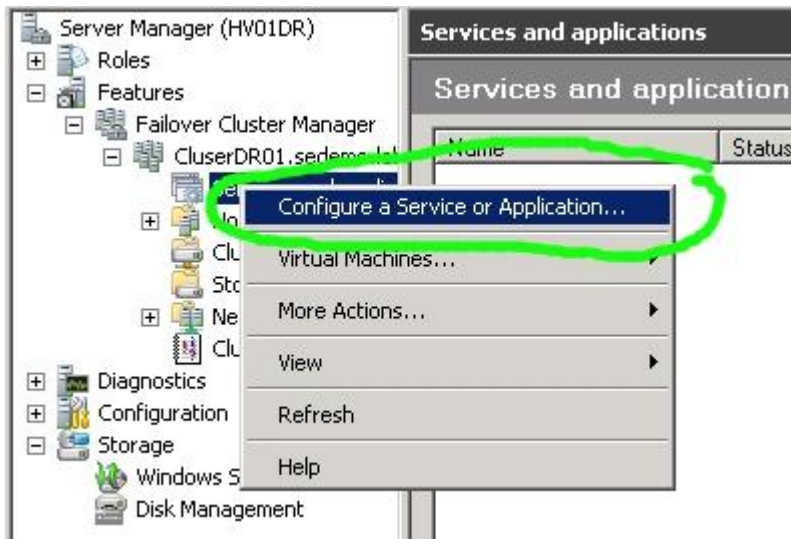
virtual machine, thus if you see multiple configuration files and you are following Nimble best practices of one virtual machine VHD per volume then you should select the most recent configuration and can safely delete the older configurations. If you are hosting multiple virtual machines from the same volume, then you will need to run the Import-VM script for each virtual machine.



- 3.2. 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.
- 3.3. 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.
- 3.4. 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:

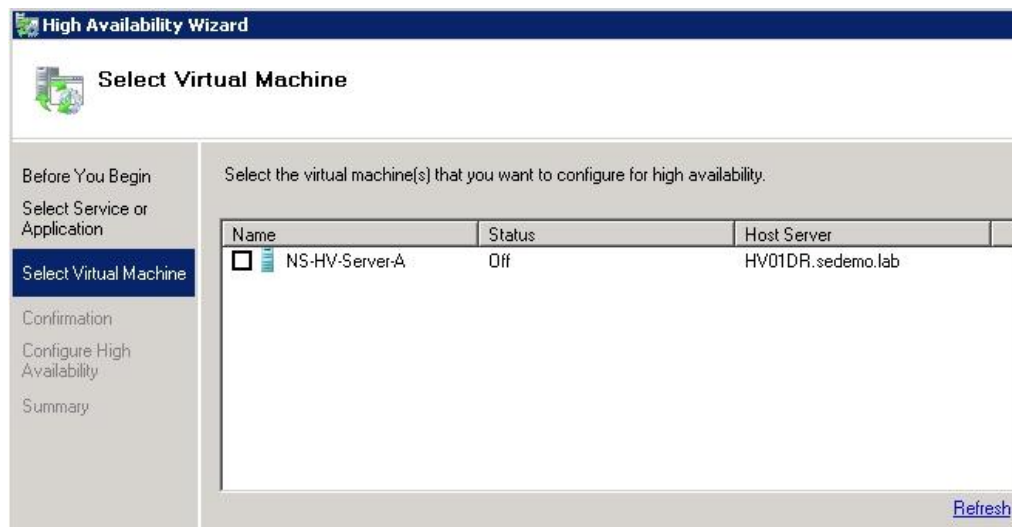
```
Import-VM 58AE37DA-53A1-412F-996E-9E26C602696D "C:\ClusterStorage\Volume1\NS-HV-Server-A"
```

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.

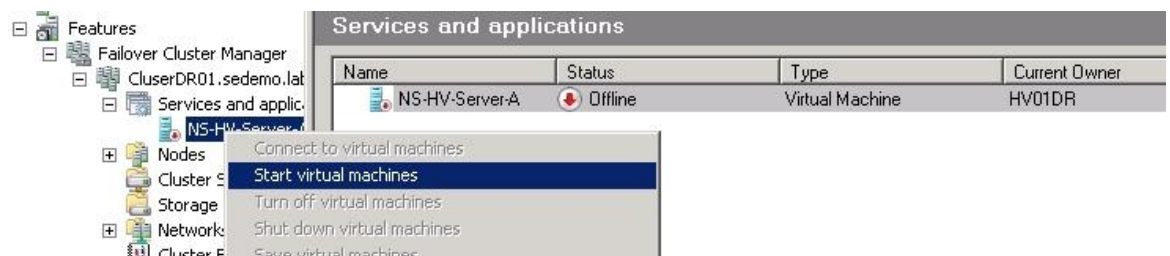


4.2. Select Virtual Machine as the resource type.

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

© 2012 Nimble Storage, Inc. All rights reserved. CASL is a trademark of Nimble Storage Inc. BPG-HPV-0812