VEEAM
IT JUST WORKS!™

Introduction to Hyper-V High-Availability

Symon Perriman
Vice-President, 5nine Software

AVAILABILITY™
for the Modern Data Center

Contents

Introduction to Hyper-V High-Availability

Businesses now need to run 24 hours a day, 7 days a week, every single day of the year. Services must be kept online or customers will be lost. In today's global marketplace, it is critical for systems to always be available so companies can remain competitive and keep their users satisfied. High data-center availability can be achieved in several ways, but the most common and easiest is through the combination of server virtualization and failover clustering.

Server virtualization provides many IT department benefits, including allowing them to consolidate multiple different systems onto a single host server, which provides higher resource utilization and density, while still providing isolation for the different workloads running within each virtual machine (VM). By separating applications within each VM, the risk of an unsecured or failed component impacting other services is reduced, and the number of components that need to be patched decreases, further reducing downtime from system reboots. Because it's easy to clone VM or deploy it from a template, this simplifies deployment, minimizes human error from misconfiguration and provides rapid scalability. Virtualization also provide mobility for that VM and allows it to be moved to different virtualization hosts with no downtime for the application running within the VM, so the host can be patched or updated. Virtualization is a critical part of any modern datacenter, so it is critical to have a highly available solution that protects and recovers from unplanned downtime, as well as simplified management of planned downtime for maintenance.

There are several different high-availability solutions for VMs. The most common are clustering, load-balancing, replication and backup. Backup can be achieved using solutions such as Veeam® Backup & Replication™, System Center Data Protection Manager or Windows Server Backup to regularly copy the VM data, then recover it if the VM virtual hard disk becomes deleted or corrupted. Replication will copy the virtual hard disk and VM configuration file to a secondary location, and also start the replicated VM if the primary VM fails. The most common solutions include Veeam Backup & Replication and Windows Server's Hyper-V Replica technology. With a load-balancing solution, such as Windows Server Network Load Balancing, several VMs are configured identically and the load-balancer distributes service requests across each of the VMs fairly evenly. This process reduces the risk of any single VM becoming overloaded. Using load-balancing is an easy way to eliminate VM downtime because VMs can be individually rotated and serviced without taking the service offline. However, load-balancing only works with identical stateless VMs, which have no shared or centralized data, such as a website front end.

Failover clustering provides the very best high-availability solution for VMs, including fast and automated recovery, access to shared data and almost no downtime due to servicing. It offers high availability for planned downtime, such as servicing, maintenance, upgrades and for unplanned downtime due to security threats, power outages or even natural disasters. This white paper will introduce the scenarios, design considerations, features and best practices for keeping services running on Microsoft's virtualization and clustering platforms, Windows Server Hyper-V and Failover Clustering online.

How Failover Clustering Works

A failover cluster is a complex distributed system, which requires interoperability between the application, operating system, VM, virtualization host, storage and networking. I'll quickly review the key requirements now and tell you how they each provide high-availability for VMs.

- **Deploys redundant hardware everywhere.** It is essential that you eliminate every single point of failure in your highly available solution. In addition to having multiple host servers (known as cluster nodes), you must have redundancy in the storage fabric by having multiple, interconnected paths to that storage, such as using MPIO (multipath **I/O**). The storage itself should use data availability features such as **RAID**, mirroring and backup. The network should have multiple paths and all network adapters should use **NIC** teaming.
- **Stores data on shared storage that's accessible by all nodes.** Since the application running on a cluster needs to share some type of information (such as a virtual hard disk for a VM or a database for a **SQL** Server), this data must be kept in a single location to which every node has access. This enables the application to run on any node in the cluster, yet still read and write the same centralized information.
- The cluster writes its configuration to the registry. Windows Server contains a built-in, hierarchical database to track system and application settings, which are known as the registry. Failover clustering leverages the registry by writing information that would be needed when the application runs on that node, specifically:
  - On which node the workload should be running.
  - The properties of each clustered VM or application such as its name, IP address, disks, networks, etc.
  - The state of each clustered VM or application such as whether it should be online, offline, failed, etc.
- **The cluster registry replicates across all nodes.** In order to ensure that a VM operates the same on every node, the settings in the cluster registry must be kept consistent across every node. This is managed by a hidden cluster-replication service that copies changes to the properties or the state of the VM to all nodes in the cluster.
- **The cluster nodes health check each other.** By using a simple request and response test, each cluster node checks the health of the cluster nodes to ensure that all the nodes are online and responsive. If a particular node misses too many health checks, the other nodes in the cluster will determine if it has failed and take corrective actions to restart any services that were running on it before it crashed.
- **Restarts clustered workloads using data from the cluster registry.** Even if one cluster node crashes, every other node in the cluster knows which VMs the failed node was hosting through the data replicated in the registry. The healthy nodes will then use this information to restart the VMs in the same state they were before the node crashed. It will configure the same properties (name, IP address, disk, state, etc.) and the same state (online, offline, etc.), and access the same data from the shared storage. This essentially recreates an identical VM with the same connection settings on the healthy cluster node. Any services or clients that were using the VM before the node crashed will then be able to reconnect to this new VM, which will appear identical to the VM that they were using before the crash.

Deploy a Failover Cluster

Since a failover cluster requires many hardware and software components, there are multiple factors that need to be considered during deployment. Most the following administrative tasks are done through the **Failover Clustering GUI** console, known as **Failover Cluster Manager**, although the tasks can also be done through PowerShell. For more information about managing a cluster with PowerShell, click here: https://technet.microsoft.com/en-us/library/hh847239.aspx.

Host SKU Selection

There are several versions of Windows Server that support failover clustering and they offer an identical feature set. There are a few differences that apply to each SKU described below:

- **Windows Server Datacenter Edition** — This SKU is recommended for virtualized environments since it allows you to run an unlimited number of VMs that use Windows Server in the guest operating system. When you purchase a license for the host, all the Windows Server VMs running on that host can share the same license and even be automatically activated by the host using the Automatic VM Activation feature.
- **Windows Server Standard Edition** — This SKU is limited by its virtualization rights, so it is not recommend for Hyper-V hosts. With each host license, you are only provided a shared license for two additional VMs running Windows Server. This SKU is generally used for servers not hosting VMs, such as a Domain Controller or SQL Server.
- **Hyper-V Server** — This SKU, designed for Hyper-V hosts, is completely free and never expires, and is designed to provide a zero cost option to help grow Hyper-V adoption. It supports all the same virtualization and high-availability features as the other SKUs. However, it has no GUI and only comes with core virtualization, storage and networking features. It does not include any other Windows Server features and roles such as AD (Active Directory), DNS or IIS. It also does not include any licenses to run Windows Server in the guest OS. Although, if you use Linux inside the guest OS, you can get an almost entirely free Hyper-V solution. Hyper-V Server is available here: aka.ms/HyperVServer.

Both **Windows Server Datacenter Edition and Standard Edition** offer a setup choice for the full installation (with a GUI) and the Server Core installation (without a GUI). Both options support all the features of Hyper-V and Failover Clustering. Server Core is recommended for Hyper-V hosts because it provides higher availability through a reduced attack and smaller servicing footprint, and it can be managed remotely using a full installation (with a GUI) of Windows Server or using Remote Server Administration Tools (aka.ms/RSAT). Server Core installations can also be managed locally using PowerShell, or a GUI can be added to it directly using 5nine Manager from 5nine Software (www.5nine.com/Manager).

Windows 8 also supports Hyper-V, but it does not include any enterprise features, including live migration or failover clustering. However a VM can be created using Windows Client, then it can be exported and imported onto a Hyper-V host and have access to all enterprise features like any other VM.

Role & Feature Installation

The Hyper-V role and the **Failover Clustering** feature must be installed on every cluster node. In Windows Server 2012 and later, you can create a cluster with up to 64 nodes and 8,000 VMs, so role and feature installation can be scripted using PowerShell, along with any other repetitive Hyper-V or Failover Clustering management task.

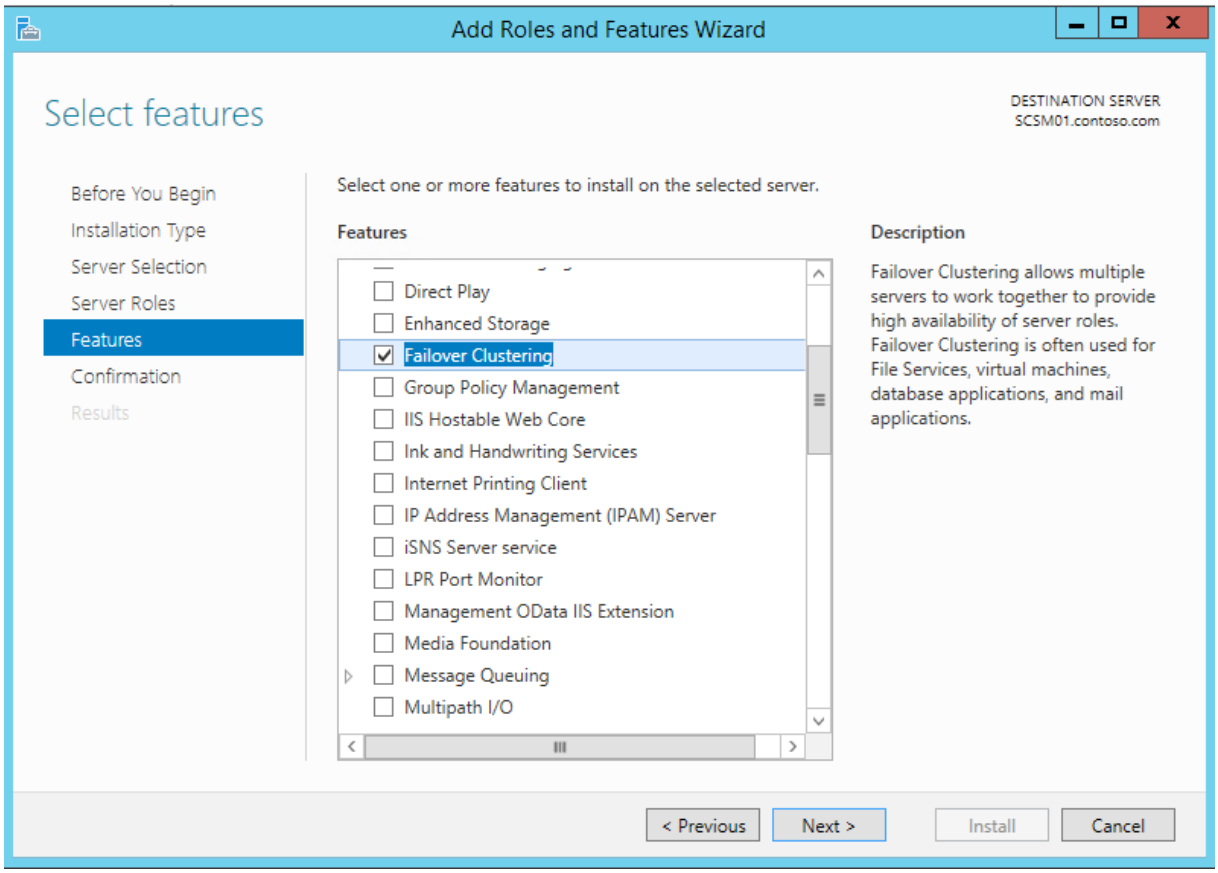

Figure 1 - Installation of the Failover Clustering Feature

Active Directory Configuration

AD is required to use **Failover Clustering**, even though Hyper-V does not require it. This means that even if you intend to use the free Hyper-V Server SKU for your cluster nodes, you still need at least one full version of Windows Server running AD in your environment. All cluster nodes must be part of the same AD domain running at a Domain Functional Level of at least Window Server 2003. The cluster nodes must also be member servers, which means that you cannot run the domain controller directly on the host. These requirements are because AD is required for authentication when starting clustered VMs and other workloads, so if the cluster node hosting the Domain Controller has its communications disrupted or blocked by cluster network traffic it may not be able to start the application. It is also a good best practice to have all the cluster node computer objects in the same Organization Unit so that Group Policy can be applied consistently across all nodes. If BitLocker is enabled to protect the data of the VMs from physical theft, then the Domain Controller must be running Windows Server 2012 R2.

It is also possible to run the Domain Controller in a dedicated clustered VM, providing it with the same high-availability benefits of clustering any VM. In addition, you can preconfigure cluster accounts in AD so that administrators with lesser privileges are able to create clusters or VMs in a Read-Only Domain Controller (RODC) environment. While this topic is beyond the scope of this white paper, more information can be found here.

Storage Configuration

Failover clustering requires some type of shared storage so the data can be accessed by all nodes and will work with most modern SANs (storage area networks) using a supported protocol. Fortunately, failover clustering includes a built-in, best-practice analytics tool (Cluster Validation) that will verify that the storage will work correctly. You can all buy a pre-tested solution from the Windows Server catalog) and the Microsoft Private Cloud Fast Track program. Many Hyper-V users that are purchasing new hardware are choosing converged systems that bundle all the cluster nodes, storage and networking components into a single rack, such as StarWind Hyper-Converged Platform, which includes Dell servers and storage provided by xByte Technologies, storage management using StarWind Virtual SAN, centralized Hyper-V and Cluster management 5nine Manager and backup using Veeam Backup & Replication.

The storage types fall into three main categories:

- **SAN using a Host Bus Adapter (HBA**) — This is the most traditional type of SAN. Supported types include Fibre Channel and Serial-Attached SCSI (SAS). Fibre Channel tends to be more expensive, but offers faster performance than SAS.
- **SAN using Ethernet** — In recent years, network bandwidth has become significantly faster, matching speeds that were previously only possible with HBA-based, storage fabric. This has enabled Ethernet-based solutions to be offered at much lower costs, though they still require dedicated NICs and networks. The two protocols supported by failover clustering are iSCSI and Fibre Channel over Ethernet (FCoE).
- **SMB 3 File Server** — The Server Message Block (SMB) protocol is a Microsoft-centric, application-layer network protocol used for file sharing on a File Server. A traditional file share, such as **\\MyShare**, is a location to store data that's accessible by multiple servers. With the introduction of Windows Server 2012, it has become possible to store the virtual hard disk for VM on this file share, which allows it to function as a very affordable shared storage type that allows all cluster nodes to access it at once.

With Windows Server 2008 R2, Failover Clustering introduced a software-defined, disk virtualization layer known as Cluster Shared Volumes (CSV), which enables a single LUN to store multiple VMs that can run on different cluster nodes. When deploying failover clustering for VMs with a SAN, it is strongly recommended to turn CSV on for all your shared disks to simplify storage management by allowing you to consolidate many VMs on a single disk. Traditional cluster disks do not permit multiple nodes to access the same disk at the same time, so you need to assign one LUN to each VM, which adds complexity to storage management.

Networking Configuration

Optimizing the cluster's networks is critical for high availability because networks are used for administration, VM access, health-checking, live migration, and often, storage if using an Ethernet-based solution or Hyper-V over SMB. The cluster nodes can be on the same subnet or different subnets, and the cluster will automatically configure the networks when the cluster is created or a new network is added.

You must use **Hyper-V Manager** to create identical virtual networks and switches on every cluster node so that your VMs can connect to other services. These virtual networks must be named the same on every node in the cluster, so the VM will always be able to connect to the same network using its name, regardless of which host the VM is running on.

All clusters require at least two networks for redundancy. If one network becomes unavailable the traffic will then be rerouted through the secondary network. The best practice is to have a dedicated network of at least 1 Gbps for each network-traffic type, which could be up to five networks:

- **Internal Cluster Traffic (required)** — In addition to using this network for health checks, the cluster needs this dedicated network to update the cluster database registry when a change happens, along with other types of internal communication. This network should not be used by other types of traffic, which could interfere with the cluster's health checking mechanism.
- **Client and Application Traffic (required)** — VMs rarely run in an isolated environment, they provide services to other application or clients. This means that those end users need to connect to the VMs through an external network. For security reasons, it is important to separate this traffic from internal cluster traffic to ensure that a denial-of-service attack on this network does not interfere with cluster traffic.
- **Live Migration Traffic (strongly recommended)** — To move a running VM from one host to another, known as a live migration, the memory of the VM is copied between the hosts through a network connection. This data causes a large spike in network traffic as several gigabytes of data are sent through the network as fast as possible. A dedicated network is strongly recommended so it doesn't interfere with other network traffic.
- **Host Management Traffic (recommended)** — There are certain types of admin tasks that also require large amounts of data to be sent through a network, such as performing a backup with Veeam Backup & Replication, deploying a VM on a host from a library or replicating a VM. Ideally, this type of traffic should have a dedicated network. If needed, however, this can be combined with the live migration traffic network.
- **Storage using Ethernet Traffic (required if using Ethernet-based storage)** — If you are using iSCSI, Fibre Channel over Ethernet or an SMB 3 File Server for storage, you must have a dedicated network connection for this storage data. This is important to ensure that this network has enough bandwidth to support the needs of all VMs on that host, or else the performance of all the VMs on a host will slow down if they cannot access data fast enough.

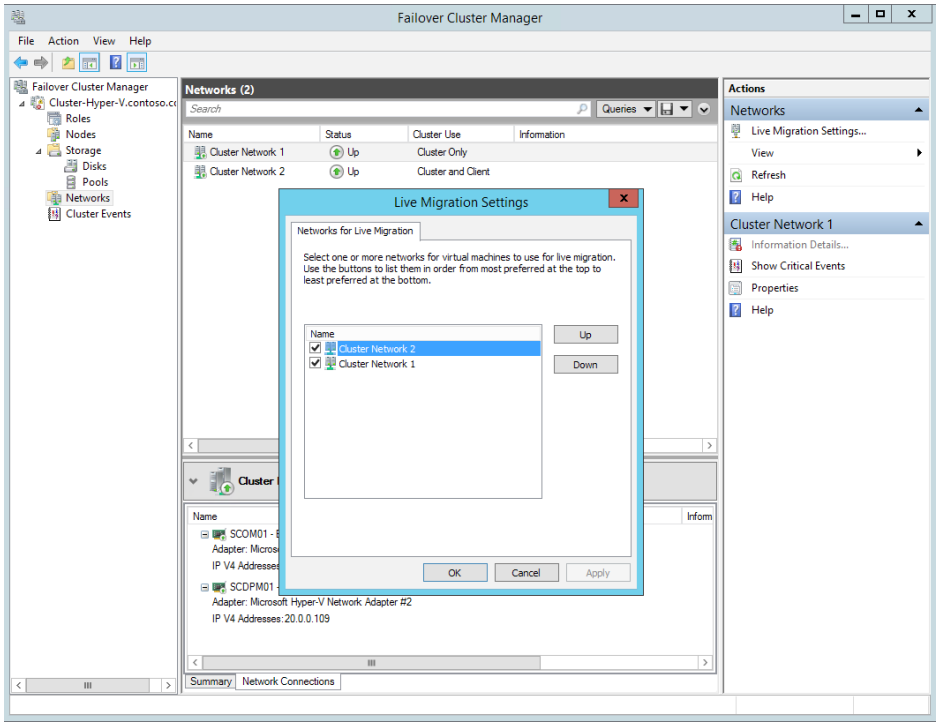

Figure 2 -- Configuration of the Live Migration Network

When the cluster is created, it will assign a different value to each of the networks based on the order that it discovers different network adapters. NICs, which have access to a default gateway, are designated for use with client and application traffic because the cluster assumes that this network has an external connection. This value is known as the Network Priority. PowerShell must be used to configure the different traffic patterns for each cluster network. Details about the network configuration can be found in this blog.

Validate a Failover Cluster

During the process of creating a failover cluster, it is possible to test every component to ensure that the solution will function correctly by using the built-in, best-practice analyzer, known as the Cluster **Validation Wizard**. This tool is launched through Failover Cluster Manager and will run a series of tests against the entire solution, including the storage, networking, Windows Server configuration and Hyper-V configuration. It will also document the settings of the cluster nodes, which can be helpful for support or compliance. This tool can even test an existing cluster as a troubleshooting tool. An additional set of tests are run to ensure that every clustered workload has been optimized for high availability.

The test result are reported in an xml-based webpage and a copy is saved on each node in the cluster under **C:\Windows\Cluster\Reports**. Each of the dozens of tests will report whether your cluster has passed or failed, or there is a warning, which means that the solution will work, but a best practice has not been followed. For a cluster to be fully supported by Microsoft, it has to meet two requirements. First, none of the cluster validation tests can fail (warnings are OK), and second, every hardware and software component must be certified for Windows Server by the manufacturer. The cluster nodes can even have different hardware, however, it is recommended that the nodes are as similar as possible to ensure that the VMs behave the same when running on any cluster node.

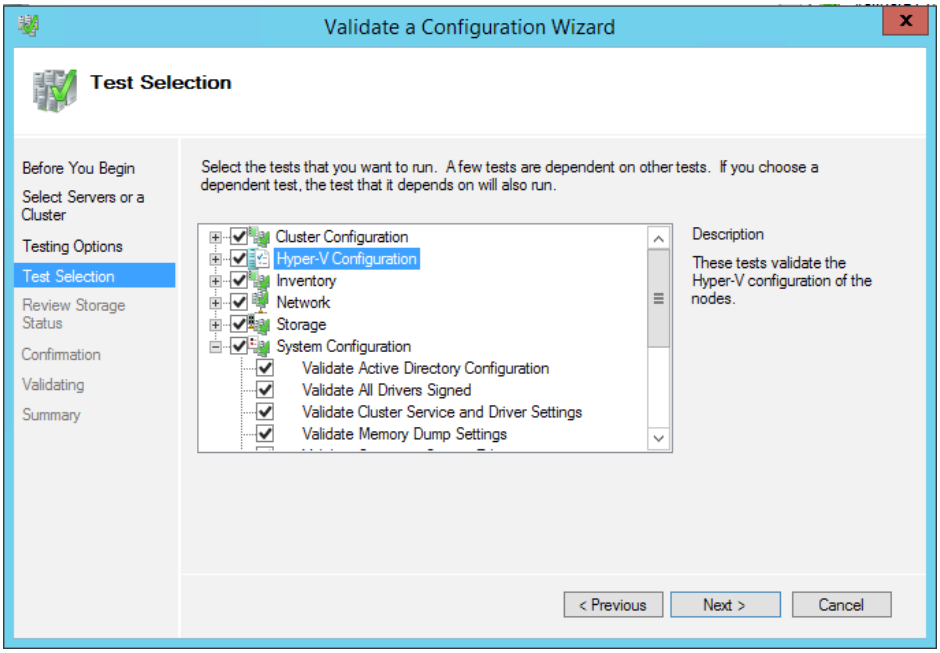

Figure 3 - Cluster Validation to Test the Configuration

Create a Failover Cluster

After the hardware and software configuration has passed cluster validation, the cluster can be created in a few steps using the **Create a Cluster Wizard**, which is launched from **Failover Cluster Manager**. After specifying the Hyper-V hosts that will become cluster nodes, you will be prompted to provide an IP address and NetBIOS name for the cluster. If a DHCP Server is detected, it will automatically provide the cluster with an IPv4 or IPv6 address and this can be changed later. This cluster name is actually another clustered resource that will run on any active node and failover to a different node if that first node becomes unavailable. This allows you to connect to the cluster by only specifying this cluster name or IP address, without actually needing to know which of the cluster nodes are currently online.

Create a Scale-Out File Server

If you decide to use a SMB 3.0 File Server to access the VM's virtual hard disk, this is provided through a file share that needs to run on a file server. If that file-share path to the virtual hard disk, such as **MyShare\MyVM.vhdx**, becomes unavailable, then the VM will not be able access its disk and it will not function correctly. For this reason, it is important to make access to this file share highly available. This can also be done with a **Failover Cluster**, although it should be done on an independent cluster because the Hyper-V cluster's hardware needs to be optimized for virtualization.

Creating a highly available File Server is done by the High Availability Wizard, which is launched through Failover Cluster Manager. However, this type of File Server can still cause temporary VM downtime if the node hosting the file share becomes unavailable and has to failover, since the file share will be temporarily offline.

In Windows Server 2012, a new type of clustered File Server was introduced called Scale-Out File Server (SOFS). This enables the File Share path to be created at the same time across several nodes simultaneously by leveraging the Cluster Shared Volumes (CSV) technology. This means that several nodes can fail, but as long as one of the nodes hosting a file path stays online, then the VMs will have continual access to their virtual hard disks and stay running. The Scale-Out File Server is not required, but it is recommended to provide the highest-levels of availability for your clustered VMs.

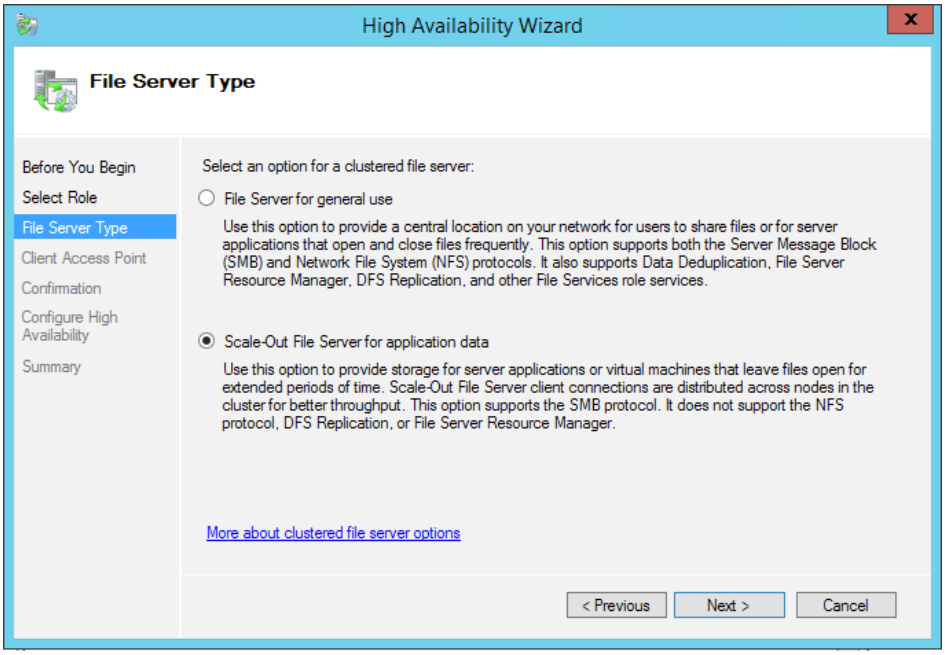

Figure 4 - Creation of a Scale-Out File Server

Deploy a Clustered Virtual Machine

Now that your failover cluster has been created, it is time to deploy your highly available VMs. There are several different ways to do this through Failover Cluster Manager, depending on whether this is a new or existing VM.

Create a Clustered Virtual Machine

Creating a VM on a cluster is slightly different from creating other clustered workloads because it does not use the **High Availability Wizard**. Instead, it uses the New Virtual Machine Wizard, identical to the tool from Hyper-V Manager. This provides a consistent experience, whether creating a clustered or standalone VM, with the main difference being that the clustered VM should be stored on shared storage accessible by every cluster node. After the New Virtual Machine Wizard has been completed, a second wizard that adds the VM to the cluster is automatically launched and completed without requiring any additional information. The VM then appears on the cluster in an offline state so you can configure additional settings before starting the VM. This uses the same properties interface as Hyper-V Manager. You can run any type of guest OS on a cluster that is supported by Hyper-V, and a current list of Windows Server, Windows, and Linux guest operating systems can be found here.

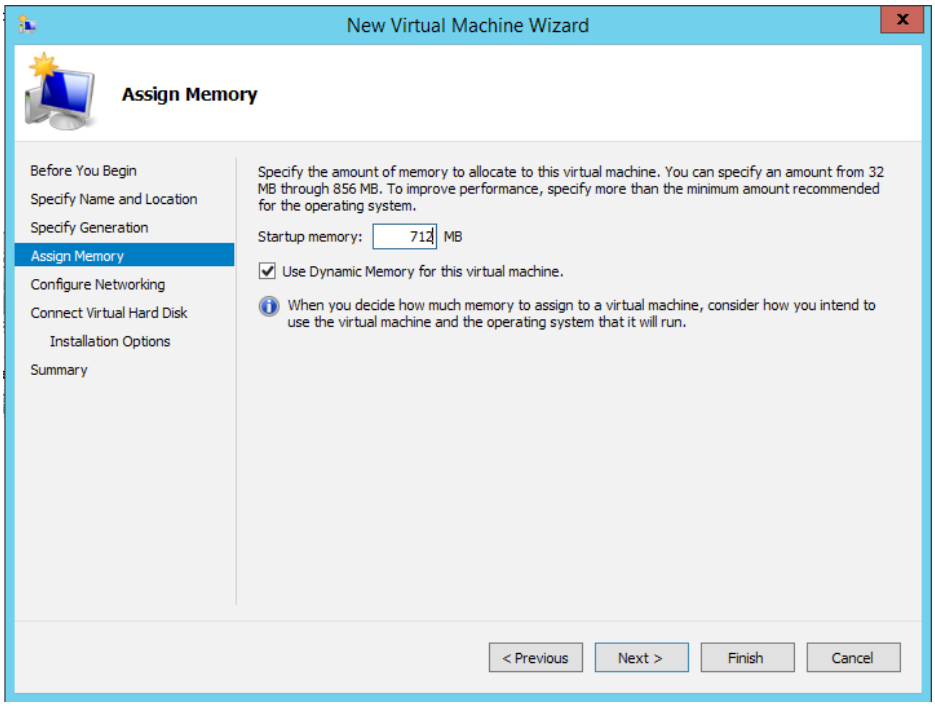

Figure 5 - Creation of a Highly Available Virtual Machine

Add an Existing Virtual Machine to a Cluster

If you wish to make an existing VM highly available, it must first be moved to one of the cluster nodes. This can be done through a live migration, an export/import or by restoring the VM from a backup. It is important to move the virtual hard disk for the VM to shared storage managed by the cluster. Once the VM is hosted on a cluster node, it can be made highly available through the High Availability Wizard to select the VM resource type. If you are using Veeam Backup & Replication, you can restore the VM to the cluster and immediately make it highly available. The VM will now be managed by Failover Cluster Manager and it will failover between nodes.

Migrate a Virtual Machine from an Existing Cluster

If you already have clustered VMs in your environment, it is possible to migrate them to your new cluster using the **Copy Cluster Roles Wizard**, which is launched through Failover Cluster Manager. After launching the wizard from your new cluster, you will specify the source cluster, then select those that you want to migrate from a list of VMs running on it. You will be given the option to remap the cluster storage and network paths. After completing the migration, you can offline the VM on the old cluster and online the newly migrated VM on the new cluster.

Create a Guest Failover Cluster

It is possible to provide even higher availability for virtualized workloads by creating a new cluster from virtual servers instead of physical servers. This is known as **guest clustering**. This enables an application running within the VM, such as SQL Server, to be highly available so it's resilient if the workload fails, the VM crashes or the guest OS needs to be updated or serviced, by moving the workload to a different virtual cluster node. Configuring a guest cluster follows the same steps as creating a physical failover cluster, with a few additional considerations:

- **Workload** – Instead of making a VM highly available, you will be clustering a different application running inside the VM's guest OS. Special configuration may be needed for this workload.
  NOTE: *It is not possible to run a VM within a VM (also known as nested virtualization).*
- **Storage** – Storage must also be shared by VMs, so all virtual cluster nodes can access the data. The VMs can only connect to storage that has a virtualized HBA or that's Ethernet-based because they cannot connect to a physical storage adapter. Hyper-V provides a virtualized Fibre Channel HBA to enable the VM to connect to this type of storage, but a virtualized Serial-Attached SCSI (SAS) adapter is not available for Hyper-V. The VM can also connect to storage using: iSCSI, Fibre Channel over Ethernet or an SMB 3.0 file share.
  - **Shared VHDX** – To simplify storage management for a guest cluster, it is possible to give all the cluster nodes access to a single LUN that has multiple VHDX (virtual hard disk) files on it. Each of these disks can function as shared storage and will store guest cluster data, such as the database for the virtualized SQL Server. Details about configuring the Shared VHDX disk are available here.

Manage a Clustered Virtual Machine

There are many management options for a failover cluster, and those that are important for Hyper-V clusters are described in this section.

Virtual Machine Groups

Each clustered workload is organized into its own group, which contains several dependent resources and move together between different cluster nodes. **A cluster group** for a VM contains the VM, the VM's configuration file and the disk on shared storage (either a CSV or traditional cluster disk). When the VM is migrated or fails over to a different node, ownership of the configuration file and disk also move to that node.

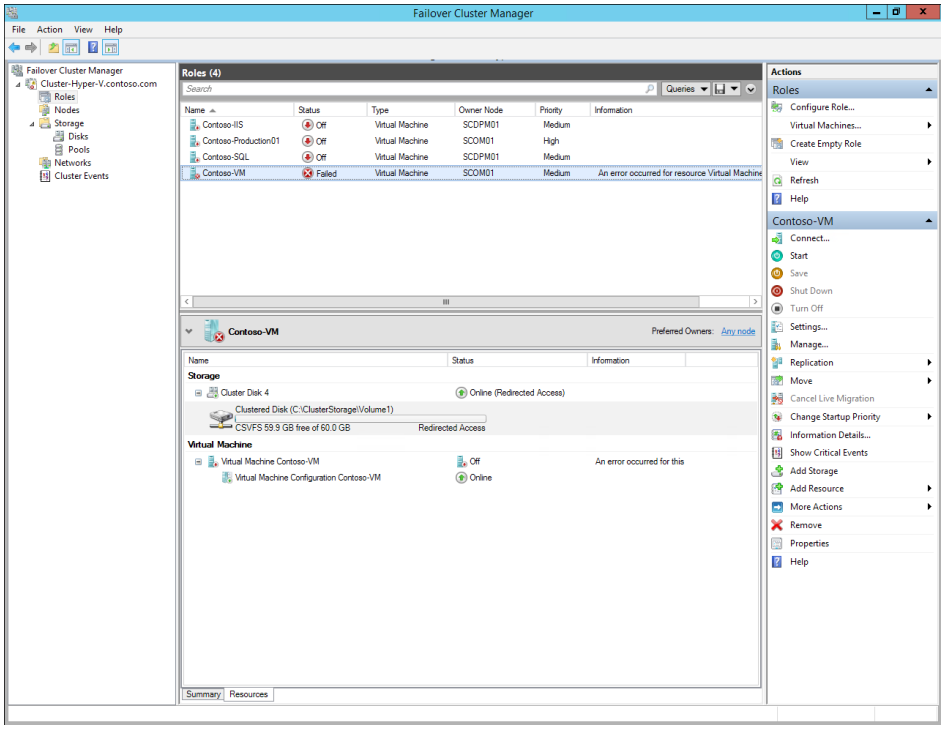

Figure 6 - Management of Several Clustered Virtual Machines

Virtual Machine Startup Priority

Clustered VMs can be assigned a startup priority setting of **High, Medium, Low** or **No Auto Start**. These four settings allow you to organize the VMs into several management groups. When an operation needs to be applied to a group of VMs, the priority setting is used to perform the action on the High VMs first, followed by Medium, Low and No Auto Start. Some operations that consider the priority include starting up VMs on a node, live migrating a group of VMs and draining VMs from a node when it is placed into maintenance mode. If your cluster crashes, this allows you to ensure that your infrastructure or business-critical VMs are restarted first without contending for resources with less important VMs.

Some admin also use the VM priority settings to organize the VMs into different tiers. For example, you could assign a High priority setting to your backend database VMs so they start first, followed by a Medium priority setting to middle-tier business VMs, and, finally, a Low priority setting to the front end web VMs so they are launched last and allow for extra time for the backend services to come online before giving clients access to the system.

The **No Auto Start** setting is slightly different because a VM assigned this priority will not automatically start if the host has crashed, which will allow the other VMs that are critical to the business to come online first. Be careful if you assign No Auto Start to a VM, as it must be manually started by the cluster administrator.

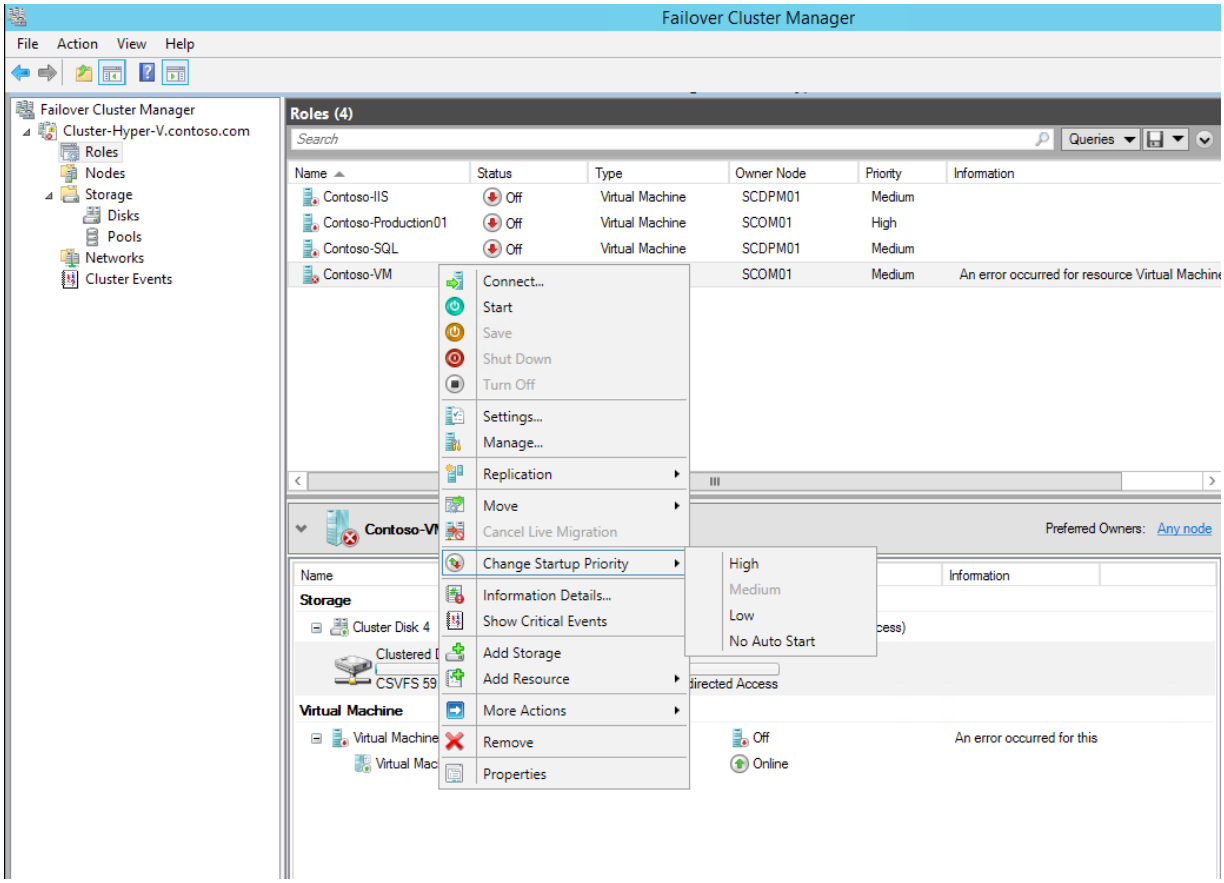

Figure 7 - Configuration of the Startup Priority for a Virtual Machine

Virtual Machine Placement on Hosts

Windows Server Failover Clustering does not require identical hardware for each host, as long as the entire solution does not fail any of the cluster validation tests. Since some hosts may be more powerful than others or have different access speeds to the storage, you may want certain VMs to run on specific hosts.

The **Preferred Owners** setting allows you to specify on which cluster node(s) you prefer the VM to run. If any of the Preferred Owners nodes are online, the VM will be hosted there. If none of the Preferred Owners nodes are online, then the VM will move to another active host. The VM can run on any cluster node and provide more failover options because, by default, none of the nodes are made Preferred Owners.

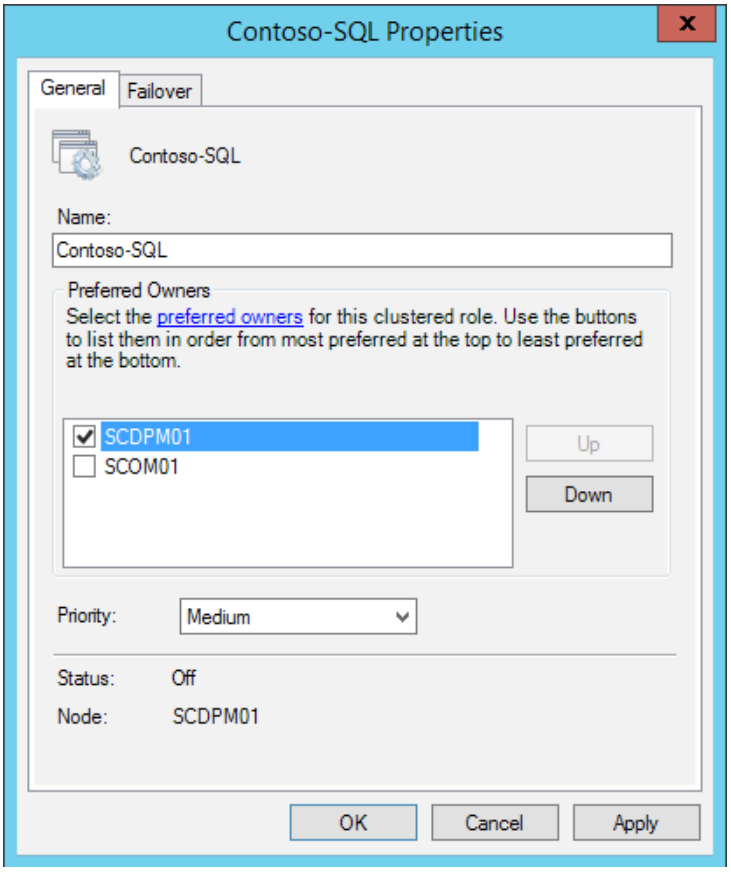

Figure 8 - Configuring Preferred Owners for a Clustered VM

The **Possible Owners** setting specifies whether the VM is allowed to run at all on a node. If none of the Possible Owners nodes are available, the VM will stay in an offline state and wait for a Possible Owners node to come back online. By default all cluster nodes are Possible Owners, meaning that the VM can run on any cluster node for higher availability.

Using PowerShell or System Center Virtual Machine Manager, you can assign VMs to be part of an **Availability Set**. All VMs in this group will deliberately try to distribute themselves across different hosts to try to provide higher availability. If you have two or more virtualized Domain Controllers, you would put those VMs into an Availability Set so that are not all running on the same host because that would create a single point of failure. Besides using this to group VMs by role, another common case is to put all nodes from the same guest (virtualized) cluster into an Availability Set. This distributes the guest cluster nodes across different physical nodes and ensures that the failure of a physical node will not take down an entire guest cluster. Using PowerShell, this is configured with the **AntiAffinityClassName** property. More information on this topic can be found in this blog.

Another host-placement setting that you should evaluate is **Failback**, which determines if you want the VMs to return to this host or remain where it is, once it comes back online after a crash. Even though VMs are moved using live migration to eliminate downtime, turning on Failback does add additional live migration network traffic, so it is disabled by default.

Virtual Machine Offline Settings

Cluster nodes will go offline occasionally due to unplanned downtime or planned maintenance. When the outage is expected, it is recommended to live migrate the VMs to a different host so there is no downtime. When there is an unplanned failure on the host, such as a power outage or blue screen, the VMs will turn off and restart on a different cluster node, or possibly on the same node, depending on the host-placement settings. Also consider whether the VM's Priority is set to No Auto Start because this would require the VM to be manually restarted if the host crashed.

When you view the properties of a VM under the **Policies** tab, you will see several settings for restarting the VM when it fails. In general, you will want to restart the VM immediately, but you may also want to add a delay if you want to review the problem or if dependent services need to be restarted first.

Under the **Settings** tab you will be able to control how a VM is shutdown when the cluster takes the resource offline. By default, it uses the **Save** option to commit the VM's running memory to the disk. You also have the option to **Shut Down**, **(Force) Shut Down**, and **Turn Off**, which either gracefully shut down the VM or immediately turn it off.

Virtual Machine Monitoring

One advantage of clustering any application is the cluster will regularly check and restart a workload through a mechanism called **Heartbeat** if it has stopped, is hung or has crashed. Clustered Hyper-V VMs use Heartbeat to ensure that the VM is running for an additional health check that verifies that the guest OS is not hung or frozen, even if the VM appears to be healthy. In Windows Server 2012 R2, both of these settings are turned on by default. In earlier versions, however, they had to be configured by editing the properties of the clustered VM.

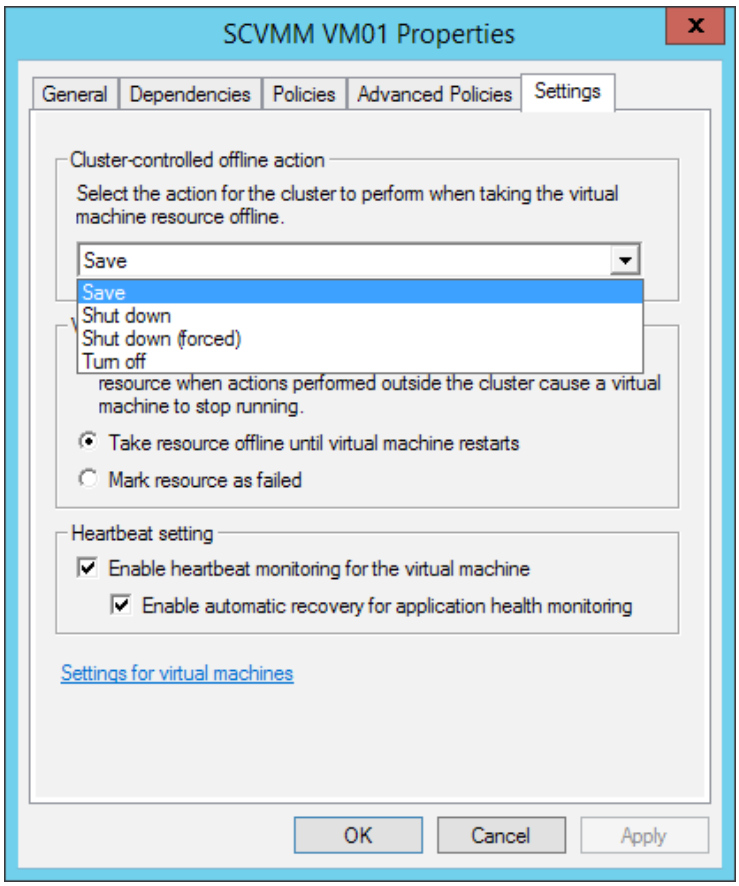

Figure 9 - Configuration of the Offline Actions & Heartbeat Settings for a Clustered VM

If you have access to the **VM's guest OS running Windows 8, Windows Server 2012 or a later version**, there is an additional setting you can configure that lets you monitor any Windows service running within the guest OS from the cluster node. This allows you to try advanced recovery actions for this faulty service, which otherwise may never have been detected. These recovery actions include restarting the service, restarting the VM, or moving the VM to a different host. Configuration for VM monitoring needs to be done on both the host and VM, and the steps are documented in this blog.

Node Health Checks

Once your VMs have been deployed, it is important to test how well they function under production conditions, especially if your networks have bandwidth constraints. The cluster uses a simple request-and-response mechanism to perform **a health check across the cluster nodes**. Slow networks could affect their reliability. These settings are configurable and are different for Hyper-V clusters than other types of clusters, which send out health checks every 5 seconds. On a Hyper-V cluster, heath checks occur every 10 seconds to nodes on the same subnet and every 20 seconds to nodes on different subnets. Every minute, a more-thorough health check is sent out. A node is considered to have failed if it does not respond to several consecutive health-check requests, which could be due to network traffic and even if the node is healthy.

The default settings should be suitable for most clusters, but, based on the bandwidth, numbers of networks, distance between nodes and network traffic, it is important to test the settings in production conditions and adjust health checks as necessary.

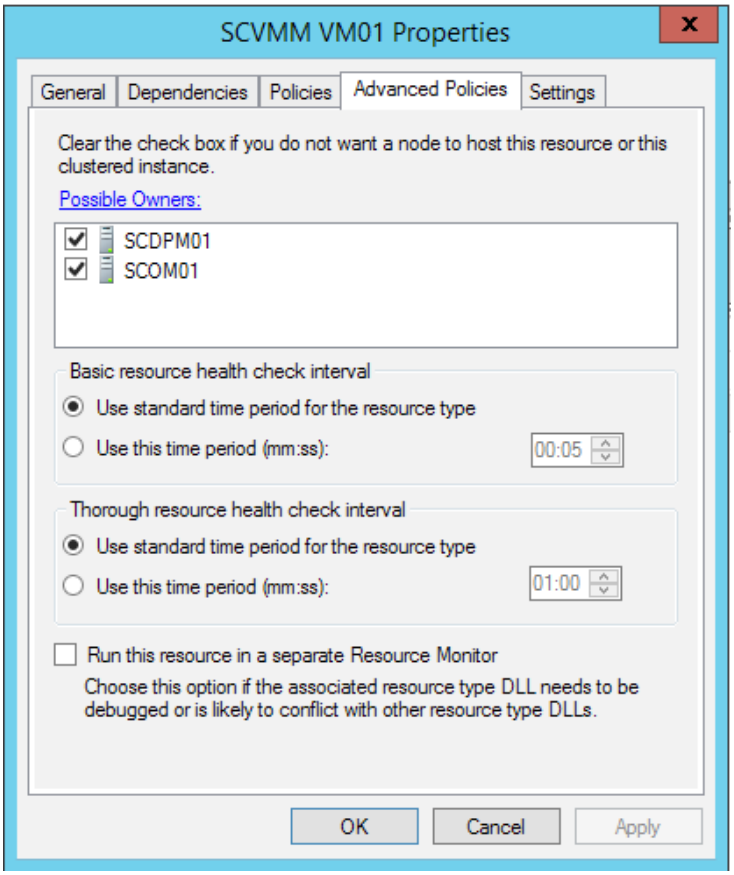

Figure 10 - Configuration of Cluster Node Health Checks

Node Maintenance

If you need to do planned maintenance for a cluster mode, there is an easy way to live migrate all of the VMs from it in a single step. By selecting a node, you can **Pause** it, and have the option to **Drain Roles** or not. When you **Pause** a node, this places the node into maintenance, which means that no VMs will failover onto it. If you choose to **Drain Roles**, the cluster will live migrate the VMs onto other cluster nodes, based on their **Priority**. To remove a node from maintenance, you select the node and Resume it, and you will have the option to **Fail Roles Back** or not, which will live migrate back all of the VMs that this node previously hosted.

There is a cluster-wide property, **DrainOnShutdown**, which is enabled by default. When an admin shuts down a Hyper-V host, the shutdown will be delayed while the VMs are live migrated from it to other nodes, starting with the **High** priority VMs. This prevents downtime if an admin tries to shut down the cluster node, but is unaware that there are VMs running on it.

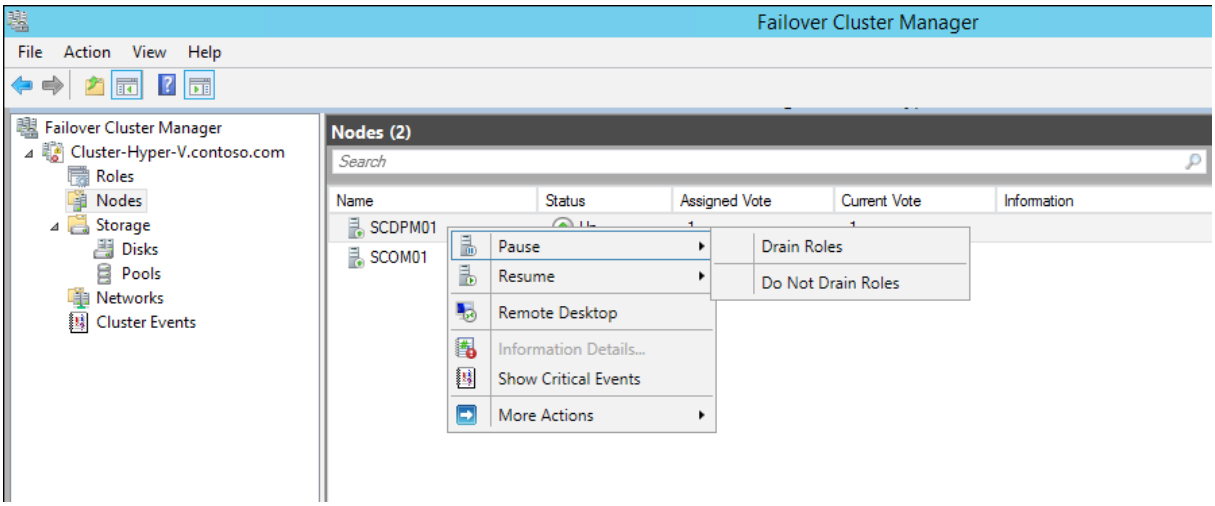

Figure 11 - Placement of a Node into Maintenance Mode

Cluster Aware Updating

It is important to keep your VMs updated with patches for your software and hardware. Windows Server 2012 introduced a feature called **Cluster Aware Updating (CAU)**, which automates the repetitive cluster patching process on both physical and guest (virtual) clusters. CAU can be run locally or remotely to centrally manage cluster patching. It connects to a Windows Server Update Services server or Windows Update. The updating workflow follows these steps:

1. The administrator creates a collection of updates, known as a **Baseline**, for the cluster.
2. CAU compares the updates on each node to the Baseline and determines which nodes need which patches.
3. CAU selects the node with the fewest workloads to minimize disruption, then **Pause** it and place it into maintenance and then live-migrate the VMs to other nodes.
4. After all VMs have been drained from the node, updates are downloaded and applied to the node and it is restarted (if necessary).
5. After the node comes back online after the restart, CAU will verify that the patches were installed successfully and the cluster node is functioning correctly.
6. CAU will **Resume** the node to take it out of maintenance and allow it to host new VMs again.
   NOTE: *CAU repeats this process for all the nodes in the cluster.*

If CAU encounters an error, it will suspend the patching cycle and alert the administrator. This is to ensure that a faulty update is not applied to every node in the cluster. CAU also has a self-updating mode that allows it to run on the same cluster it is patching. This allows the updating service itself to move to different cluster nodes as they are restarted, which removes the need to have any dedicated patching infrastructure.

Clustered Virtual Machine Security

Securing virtualized environments from viruses, malware and other threats is different than in traditional data centers where the server hardware is generally static. In virtualized environments, it is easy to install and manage agents on every server through **endpoint protection**. Virtualized environments need to manage security differently than traditional data centers because virtual machines, disks and networks can be dynamic and constantly changing. It's usually impractical to install security agents inside every VM. **5nine Cloud Security** provides the only agentless antivirus, firewall and intrusion detection solution for Hyper-V. This security software filters traffic going into and out of the VMs through an extension to the virtual switch, which provides protection at the host level and before the threat even reaches the VM. This means security is centrally managed and the VM user never has to worry about updating or scanning the guest OS, regardless of whether it is running Windows Server, Windows, or Linux. Get more information about 5nine Cloud Security.

Clustered Virtual Machine Replication

Windows Server 2012 introduced a new feature for Hyper-V that provides Hyper-V VM disaster recovery by replicating virtual hard disks to a different datacenter, or even to Microsoft Azure. The **Hyper-V Replica** feature needs to be enabled for any server or cluster node that's hosting or receiving a replicated VM. Each VM can be individually configured and tested for replication, and then send a copy of the data asynchronously every 30 seconds, 5 minutes or 15 minutes.

The Hyper-V Replica on a failover cluster is configured differently by creating a new clustered workload called the **Hyper-V Replica Broker**. This is a highly available, replication-service version, which means that the replication engine will failover between cluster nodes to ensure that replication always happens.

NOTE: *If you are configuring replication on a cluster, be sure to store the virtual hard disks on shared storage, rather than to a local disk.*

Designing a DR plan for a cluster should be based on: the workload, the distance between datacenters and other network factors. This article describes the different design options with a failover cluster.

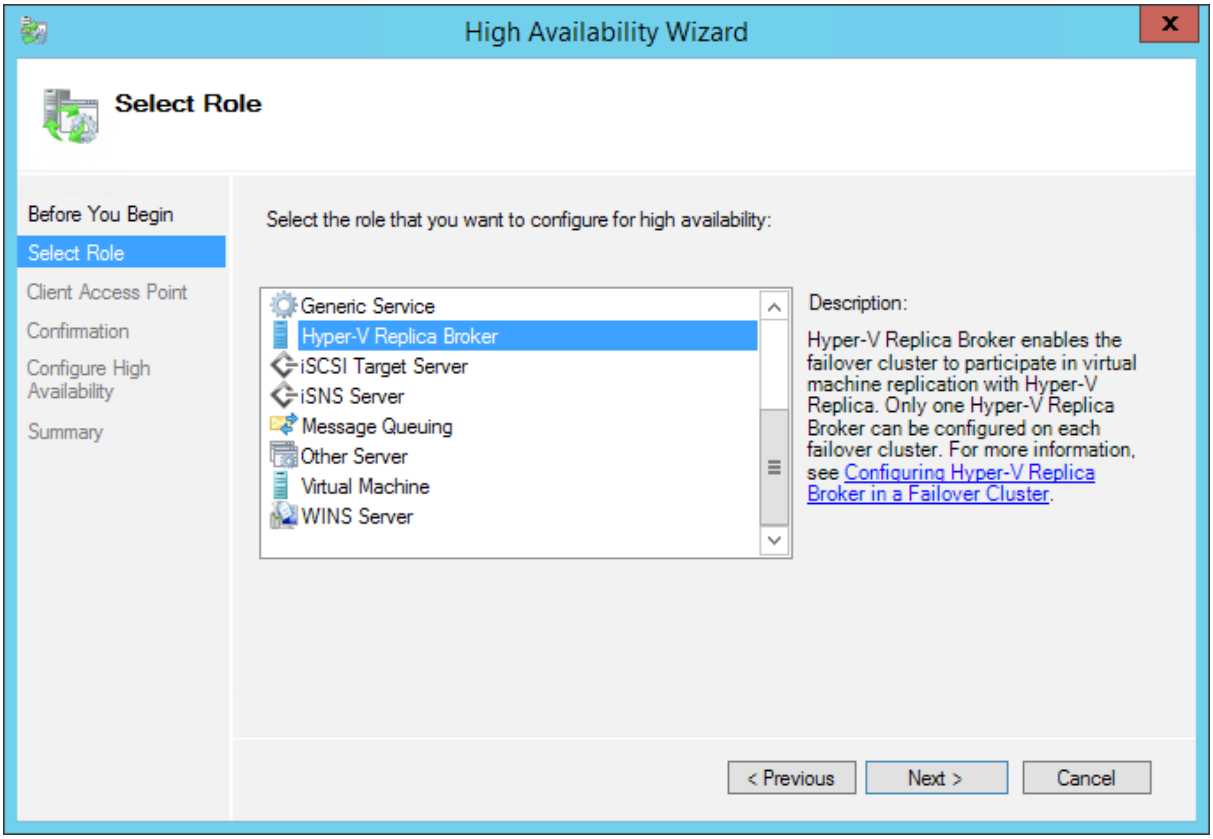

Figure 12 - Creation of the Hyper-V Replica Broker

Clustered Virtual Machine Backup

Once your clustered VMs are running, it is important to configure automatic regular backups of VM configuration files and virtual hard disks, as well as regularly test recovery. **Veeam Backup & Replication** supports backups from all types of supported Hyper-V Clustering technologies. Whether Clustered Shared Volumes (CSVs) or SMB 3.0 shares are used, Veeam can protect these Hyper-V VMs with ease. Plus, if System Center Virtual Machine Manager is in use, backup and replication jobs will be supported.

NOTE: *System Center Virtual Machine Manager is not required with Veeam Backup & Replication.*

Virtual Machine Mobility

A major server virtualization benefit is the ability to easily move the VMs between different Hyper-V hosts with little or no downtime. It is important to remember that if you want the VM to move while maintaining service availability, the memory must be copied between the hosts and they must remain online throughout the transfer. For this reason, if a host has an unplanned failure, there will be downtime because the VM will temporarily go offline and restart on another node once the health check detects the failure.

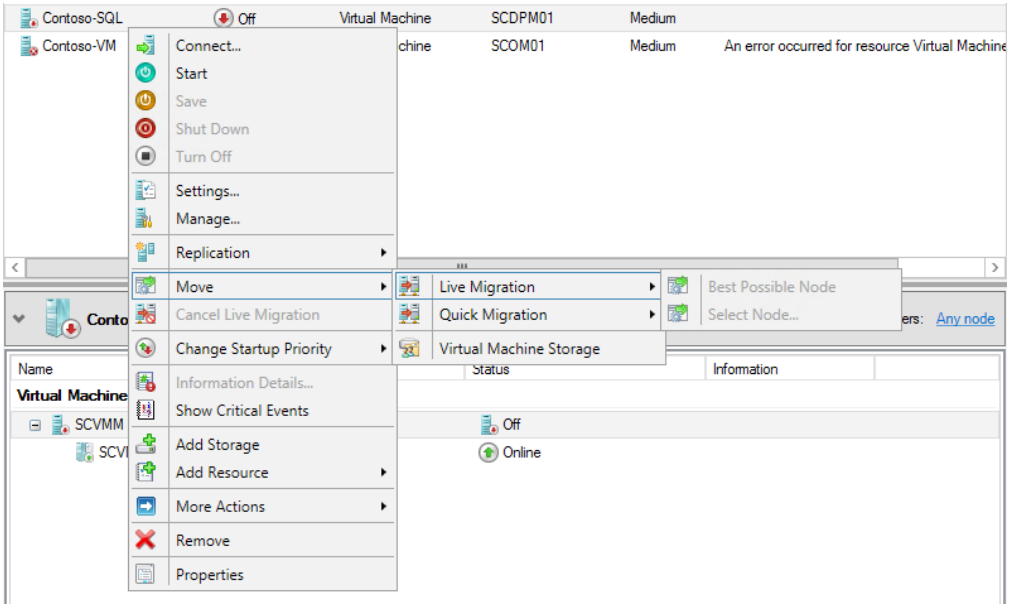

Figure 13 - Migration of a Virtual Machine to another Node

Quick Migration

The fastest way to move a VM to a different cluster node is through a **Quick Migration**, which saves the VM's running memory state to the disk, restarts the VM and loads the memory onto a new node to recreate the running VM on a new host. This method will cause downtime for the VM, so it is less commonly used, other than if you need to move all VMs off a host as quickly as possible.

When performing a **Quick Migration**, you can select which specific node to move the VM or you can also select **Best Possible Node**. This option will move the VM to the node with the largest amount of free memory, which simplifies administration since this is usually the most limiting host resource.

Live Migration

Most of the time, you will perform a live migration to move a virtual machine to other nodes or Hyper-V hosts outside the cluster. To perform a live migration, both the original and new host must be online and they should send data through the **Live Migration Network**. A live migration recursively copies the VM memory between hosts and results in something that looks like a VM clone, but on the new host. Once both VMs have consistent memory, the cluster will update network routing tables so new incoming network packets get redirected to the new VM. Finally, the old VM is destroyed so the memory it was using on the old host can be reclaimed. A live migration can be performed between cluster nodes, standalone Hyper-V hosts, or any combination of nodes and hosts.

It's possible to select multiple VMs and live migrate them simultaneously from within **Failover Cluster Manager**, and the VMs will be queued up and migrated based on priority. You can also configure a maximum number of simultaneous live migrations for each host, usually around 4 to 8, depending on your network bandwidth and workload. You can actually slow down the total live migration speed if you select a value too large, because the processors will spend too many cycles managing the recursive live migration process.

Live migration was made faster In Windows Server 2012 R2 by compressing the data before sending it over the network and then decompressing it when it has been received by the new cluster node. This feature is enabled by default because when less network traffic has to be sent between the hosts, the transfer speed is usually several times faster. Some newer networking cards support a technology called **Remote Direct Memory Access (RMDA)**, which allows even faster data-transfer speeds, making live migration up to 10x faster than standard live migration.

Storage Migration

Hyper-V also supports moving the storage location of the virtual hard disk of a VM, with no downtime for the running VM. In a process similar to live migration, Hyper-V copies the data on the virtual hard disk to a second location, then redirects future IO to that newly created disk. Using storage migration, you can perform SAN maintenance by draining the storage array—without taking your VMs offline.

There may be instances where you need to move both a running VM and its virtual hard disk at the same time. Using a **shared nothing migration**, both the VM and its disks will be copied to a new host as a single operation, with no VM downtime. This allows you to migrate a VM between a cluster and a standalone host, or even between two standalone hosts, with only local storage.

Live migration and storage migration can be enabled and configured through Hyper-V Manager by adjusting the Hyper-V **Settings**.

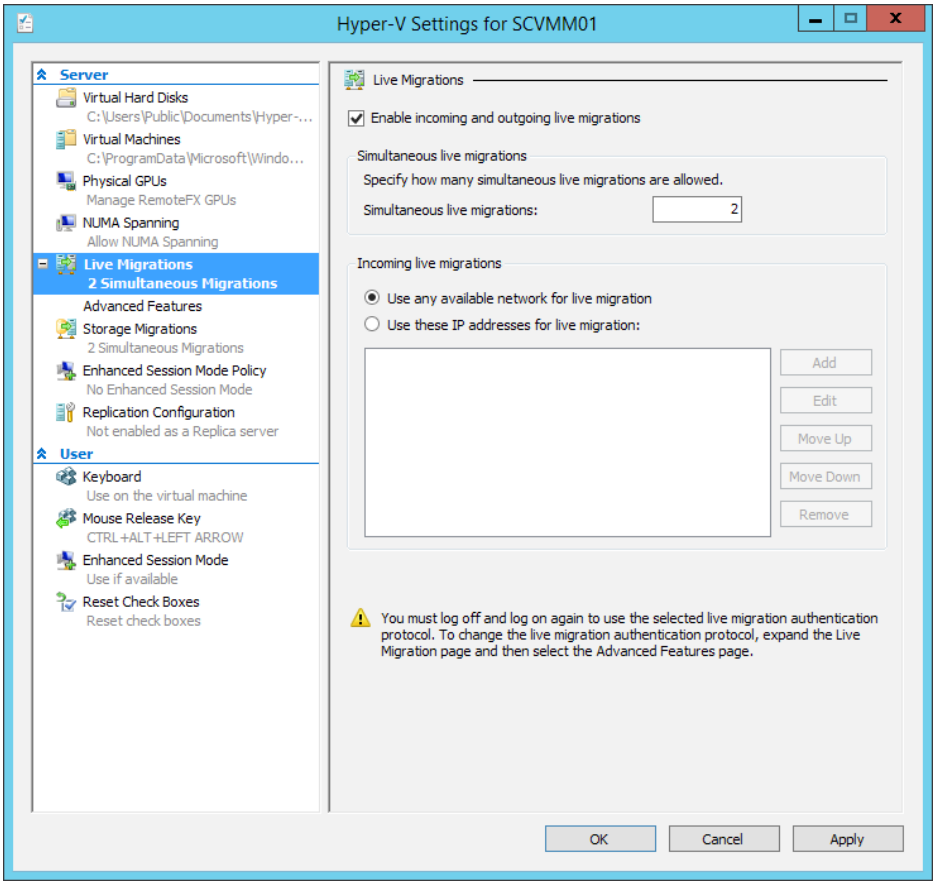

Figure 14 - Configuration of Simultaneous Live Migrations

Conclusion

Windows Server Failover Clustering should be a key piece of your data center's high availability plan because it enables continual service availability. This white paper gave you best practices for optimizing your clustered VMs. You also learned that the virtualization stack is just one part of your data center, and that it's important to eliminate any single points of failure throughout your entire data center.

Other key tips to remember include:

Ensure that your hardware is redundant and fault tolerant.

Have a disaster recovery plan in place to run your services in another data center or Microsoft Azure.

All of your management and monitoring systems need to be highly-available and as automated as possible to reduce the amount of human intervention needed to detect and resolve any issue.

Keep in mind that Microsoft's System Center 2012 R2 Management Suite provides centralized virtualization and cloud management thorough integration with Hyper-V and Failover Clustering.* System Center provides cluster management with Virtual Machine Manager, cluster monitoring with Operations Manager, automation through Orchestrator (via Virtual Machine Manager), backup through Data Protection Manager, and self-service of clustered virtual machines through Azure Pack.

*More information about integrating Windows Server Failover Clustering with System Center can be found here.

Resources

- Channel 9: Hyper-V Best Practices for High-Availability with Failover Clustering: https://channel9.msdn.com/Events/TechEd/Europe/2014/CDP-B335
- Cluster Team Blog: http://blogs.msdn.com/b/clustering/
- Microsoft Virtual Academy: Failover Clustering in Windows Server 2012 R2: https://www.microsoftvirtualacademy.com/en-US/training-courses/failover-clustering-in-windows-server-2012-r2-8489
- TechNet Failover Clustering: https://technet.microsoft.com/en-us/library/hh831579.aspx
- TechNet Hyper-V: https://technet.microsoft.com/en-us/library/mt169373.aspx
- TechNet Virtual Lab: Windows Server 2012 R2: Introduction to Failover Clustering: http://go.microsoft.com/?linkid=9846211

About the Author

Symon Perriman is 5nine Software's VP of Business Development & Marketing. Previously he was Microsoft's Senior Technical Evangelist and worldwide technical lead covering Virtualization (Hyper-V), Infrastructure (Windows Server), Management (System Center) and Cloud (Microsoft Azure). As one of Microsoft's most recognized faces, Symon has trained millions of IT Professionals, holds several patents and dozens of industry certifications, and in 2013 he co-authored "Introduction to System Center 2012 R2 for IT Professionals" (Microsoft Press).

About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24/7/365. To address this, Veeam has pioneered a new market of *Availability for the Modern Data Center*™ by helping organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. Veeam Availability Suite™, which includes Veeam Backup & Replication™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 30,500 ProPartners and more than 145,500 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit http://www.veeam.com.

VEEAM
IT JUST WORKS!™

AVAILABILITY™
for the Modern Data Center

COMING SOON

NEW Veeam® Availability Suite™ v9

RTPO™ <15 minutes for ALL applications and data
Enabling the Always-On Business™
with *Availability for the Modern Data Center*™

To learn more, visit **www.veeam.com**