



Microsoft Certification Exam 74-409: Server Virtualization with Windows Server Hyper-V and System Center

Study Guide
by Orin Thomas

Contents

Chapter 1: Virtual Machine Settings	7
Configure dynamic memory.....	7
Configure smart paging	9
Configure resource metering	11
Configure guest integration services	11
Create and configure generation 1 and generation 2 virtual machines	13
Configure and use Enhanced Session Mode	14
Configure RemoteFX	16
Summary	16
Review.....	17
Chapter 2: Virtual Machine Storage	18
Creating virtual hard disks in VHD and VHDX format	19
Configuring differencing drives.....	21
Modifying virtual hard disks.....	22
Configuring pass-through disks.....	24
Managing checkpoints	26
Implementing virtual Fibre Channel adapters.....	27
Configuring Storage Quality of Service	28
Configuring Hyper-V host clustered storage.....	29
Configuring guest cluster storage.....	30
Planning for storage optimization.....	32
Deduplication	32
Storage tiering	34
Review.....	35

Chapter 3: Hyper-V Virtual Networks and Virtualization Networking . . .	36
Hyper-V virtual switches	36
Optimizing network performance	39
Virtual machine MAC addresses	42
Configuring network isolation	44
Virtual machine network adapters	45
Virtual machine NIC teaming	47
VMM logical networks	49
VMM port profiles and logical switches	50
Network Virtualization	52
VMM virtual machine networks	53
VMM MAC and IP address pools	53
Windows Server Gateway	55
Private Virtual Local Area Networks	56
Review	57
Chapter 4: Implementing Virtual Machines	58
Highly available virtual machines	58
Guest resource optimization	61
Intelligent placement	61
Dynamic optimization	62
Power optimization	63
Performance and Resource Optimization	64
Placement rules	65
VMM templates	66
Guest OS profile	67
Hardware profile	70
VMM template configuration	74
Review	75

Chapter 5: Managing Virtualization Hosts and Infrastructure	76
Delegating virtualization management tasks	76
Role Profiles.	77
Role members.	78
Run As accounts.	79
Scope	79
Quotas	80
Networks and resources	81
Permissions	81
Managing VMM libraries.	83
Equivalent objects	84
Host group libraries.	84
Integrate third-party virtualization platforms	85
Bare metal Hyper-V host deployment	86
Integrating Operations Manager with VMM	87
Integrating Service Manager with VMM	91
Servicing virtual machine images	92
Dism and Add-WindowsPackage	92
Virtual Machine Servicing Tool 2012	93
Orchestrated Offline VM Patching Runbook.	94
Integrating Data Protection Manager.	94
Review.	98

Chapter 6: Hyper-V Failover Clustering and Failover Clustering Roles . . .	99
Cluster shared storage.	99
Cluster quorum.	100
Node Majority.	101
Node and Disk Majority	101
Node and File Share Majority	102
No Majority: Disk Only	102
Cluster node weight.	102
Dynamic quorum	103
Cluster networking.	104
Force Quorum Resiliency.	105
Cluster Aware Updating	105
Upgrading clusters.	107
Cluster Shared Volumes.	108
Active Directory detached clusters	109
Scale-out file servers	110
Preferred owner and failover settings.	110
Guest clustering	112
Shared virtual hard disk.	113
Review.	115
Chapter 7: Virtual Machine Movement.	116
Live migration	116
Storage migration.	119
Exporting, importing and copying VMs.	121
VM Network Health Detection.	122
VM drain on shutdown	124
P2V migrations	124
V2V migrations	125
Review.	125

Chapter 8: Monitoring and Hyper-V Replica	126
Audit Collection Services	127
System Center Global Service Monitor	128
Fabric monitoring	128
Operations Manager reporting	129
Operations Manager management packs	131
Monitoring Active Directory	134
Domain controller cloning	136
Hyper-V replica	137
Configuring Hyper-V Replica Servers	138
Configure VM replicas	139
Replica failover	141
Hyper-V replica broker	142
Hyper-V Recovery Manager	142
Review	143
Appendix	144
About the Author	149
About Veeam Software	149

Chapter 1: Virtual Machine Settings

The first thing you need to grasp when studying for the 74-409 exam is the basics of virtual machine (VM) settings. You'll need to understand how dynamic memory and smart paging work when VMs are starting, restarting and operating. You'll need to understand the difference between generation 1 and generation 2 VMs. You'll need to know the conditions under which Enhanced Session Mode and RemoteFX can be used. You'll also need to understand resource metering and VM guest integration services.

In this chapter you'll learn about:

- Dynamic memory
- Smart paging
- Resource metering
- Guest integration services
- Generation 1 and generation 2 VMs
- Enhanced Session Mode
- RemoteFX

Configure dynamic memory

You have two options when assigning RAM to VMs—you can assign a static amount of memory or configure dynamic memory. When you assign a static amount of memory, the amount of RAM assigned to the VM remains the same, whether the VM is starting up, running or in the process of shutting down.

When you configure dynamic memory, you configure the following values (shown in Figure 1a):

- **Startup RAM.** This is the amount of RAM allocated to the VM during startup. This can be the same as the minimum amount of RAM, or a figure up to the maximum amount of RAM allocated. Once the VM has started, it will instead use the amount of RAM configured as the Minimum RAM.
- **Minimum RAM.** This is the minimum amount of RAM that the VM will be assigned by the virtualization host. When multiple VMs are demanding memory, Hyper-V may reallocate RAM away from the VM until this Minimum RAM value is met. You can reduce the Minimum RAM setting while the VM is running, but you cannot increase it while the VM is running.

- **Maximum RAM.** This is the maximum amount of RAM that the VM will be allocated by the virtualization host. You can increase the Maximum RAM setting while the VM is running, but you cannot decrease it while the VM is running.
- **Memory buffer.** This is the percentage of memory that Hyper-V should allocate to the VM as a buffer.
- **Memory weight.** This allows you to configure how memory should be allocated to this particular VM compared to other VMs running on the same virtualization host.

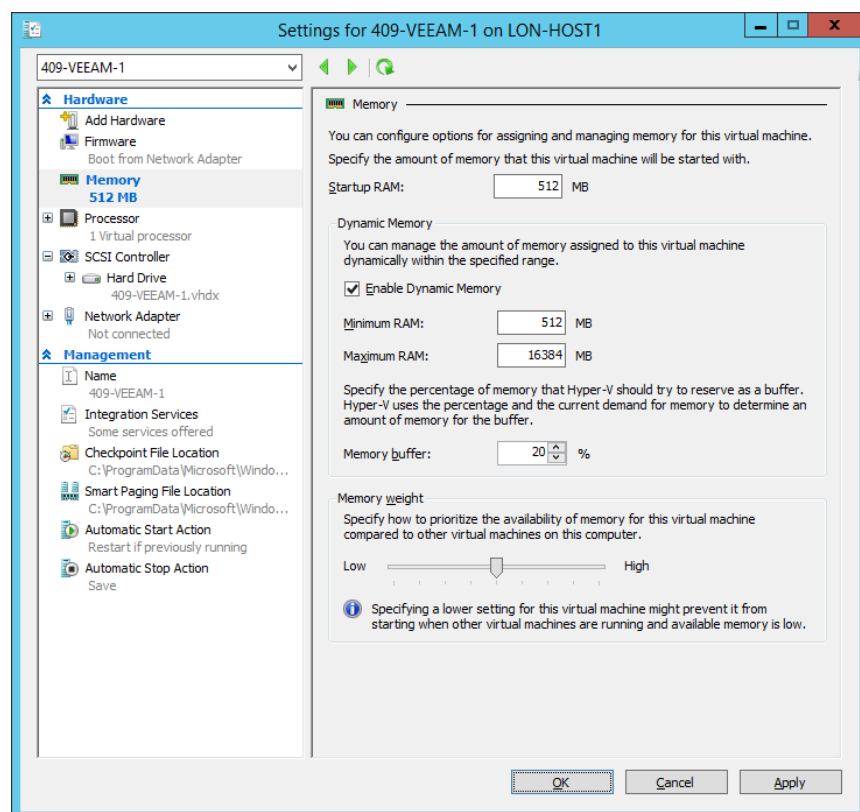


Figure 1a. Smart paging can only be used where Startup RAM exceeds Minimum RAM

Generally, when you configure dynamic memory, the amount of RAM used by a VM will fluctuate between the Minimum and Maximum RAM values. You should monitor VM RAM utilization and tune these values so that they accurately represent the VM's actual requirements. If you allocate a Minimum RAM value below what the VM would actually need to run, it is likely at some point that a shortage of RAM might cause the virtualization host to reduce the amount of memory allocated to this minimum value, causing the VM to stop running.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/library/hh831766.aspx>

Configure smart paging

Smart paging is a special technology in Hyper-V that functions in certain conditions when a VM is restarting. Smart paging uses a file on the disk to simulate memory to meet Startup RAM requirements when the Startup RAM setting exceeds the Minimum RAM setting. Startup RAM is the amount of RAM allocated to the VM when it starts, but not when it is in a running state. For example, in Figure 1b you will notice that the Startup RAM is set to 2048 MB and the Minimum RAM is set to 512 MB. In a scenario where 1024 MB of free RAM was available on the virtualization host, smart paging would allow the VM to access the 2048 MB required.

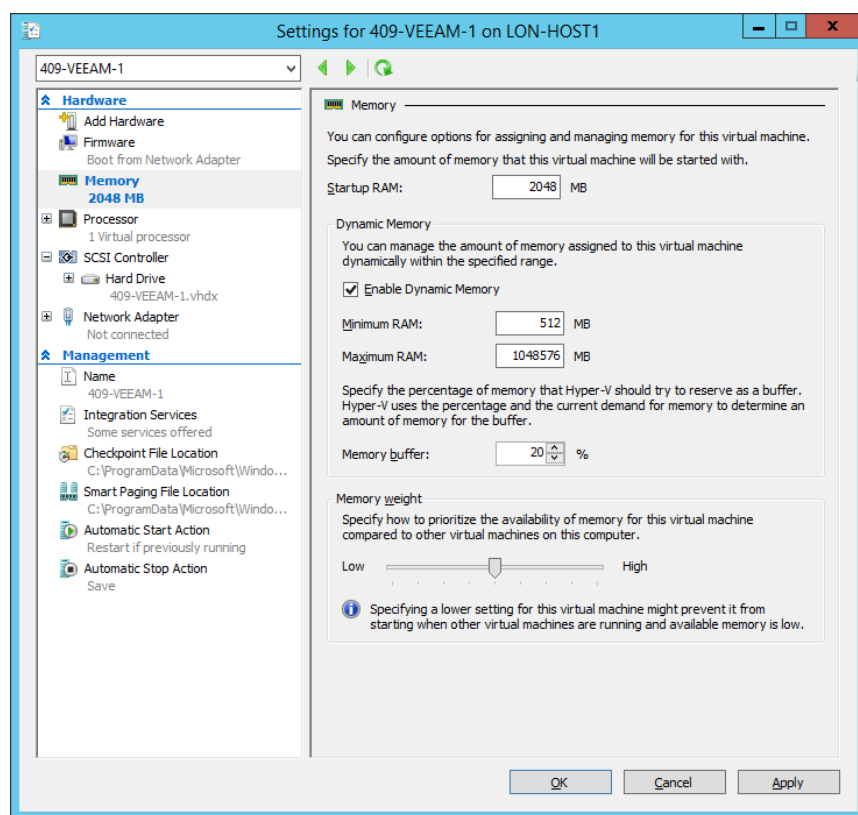


Figure 1b. Smart paging can only be used where Startup RAM exceeds Minimum RAM

Because it uses disk to simulate memory, smart paging is only active if the following three conditions occur at the same time:

- The VM is being restarted
- There is not enough memory on the virtualization host to meet the Startup RAM setting
- Memory cannot be reclaimed from other VMs running on the same host

Smart paging doesn't allow a VM to perform a "cold start" if the required amount of Startup RAM is not available but the Minimum RAM amount is. Smart paging is only used when a VM that was already running restarts and the conditions outlined earlier have been met.

You can configure the location of the smart paging file on a per-VM basis as shown in Figure 1c. By default, smart paging files are written to the C:\ProgramData\Microsoft\Windows\Hyper-V folder. The smart paging file is created only when needed and is deleted within 10 minutes of the VM restarting.




Figure 1c. Smart paging file location

Find Out More:

You can learn more about smart paging by consulting the following WindowsITPro article: <http://windowsitpro.com/hyper-v/q-i-notice-windows-server-2012-virtual-machines-have-smart-paging-file-location-what-smart-p>

Configure resource metering

Resource metering allows you to track the consumption of processor, disk, memory and network resources by individual VMs. To enable resource metering, use the **Enable-VMResourceMetering** Windows PowerShell cmdlet. You can view metering data using the **Measure-VM** Windows PowerShell cmdlet. Resource metering allows you to record the following information:

- Average CPU use
- Average memory use
- Minimum memory use
- Maximum memory use
- Maximum disk allocation
- Incoming network traffic
- Outgoing network traffic

Average CPU use is measured in megahertz (MHz). All other metrics are measured in megabytes. Although you can extract data using the **Measure-VM** cmdlet, you need to use another solution to output this data into a visual form like a graph.

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/virtualization/archive/2012/08/16/introduction-to-resource-metering.aspx>

Configure guest integration services

Integration services allow the virtualization host to extract information and perform operations on a hosted VM. It is usually necessary to install Hyper-V integration services on a VM, though Windows Server 2012 R2 and Windows 8.1 include Hyper-V integration services by default. Integration services installation files are available for all operating systems that are supported on Hyper-V. As shown in Figure 1d, you can enable the following integration services:

- **Operating system shutdown.** This integration service allows you to shut down the VM from the virtualization host, rather than from within the VM's OS.
- **Time synchronization.** Synchronizes the virtualization host's clock with the VM's clock. Ensures that the VM clock doesn't drift when the VM is started, stopped or reverted to a checkpoint.

- **Data Exchange.** Allows the virtualization host to read and modify specific VM registry values.
- **Heartbeat.** Allows the virtualization host to verify that the VM OS is still functioning and responding to requests.
- **Backup (volume checkpoint).** For VMs that support Volume Shadow Copy, this service synchronizes with the virtualization host, allowing backups of the VM while the VM is in operation.
- **Guest services.** Guest services allow you to copy files from the virtualization host to the VM using the **Copy-VMFile** Windows PowerShell cmdlet.

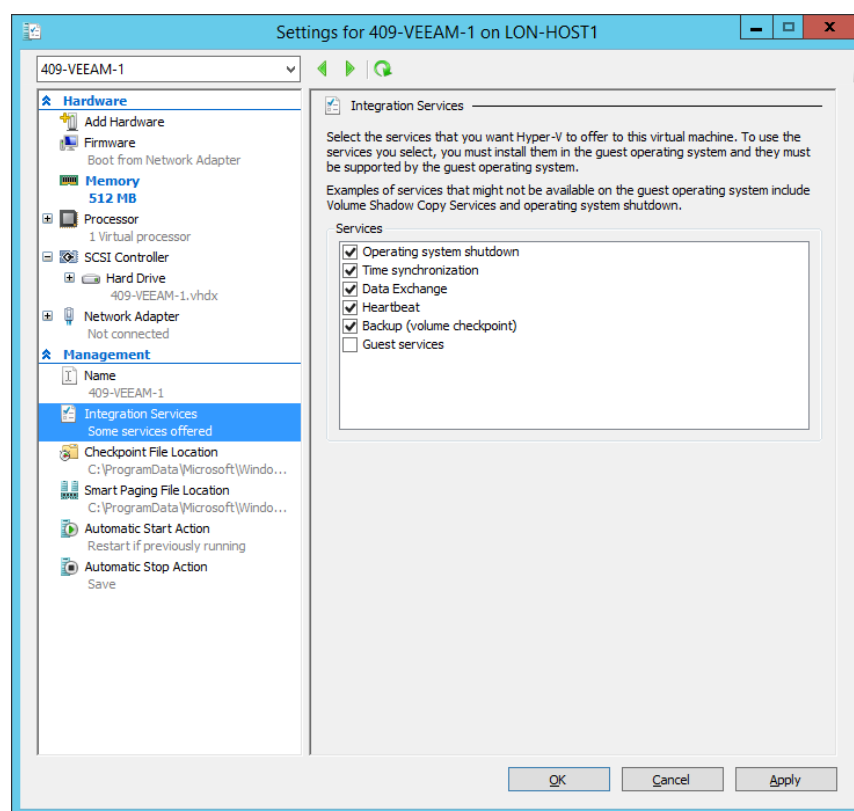


Figure 1d. Integration services

Find Out More:

You can learn more about this topic by consulting the following WindowsITPro article: <http://windowsitpro.com/hyper-v/what-guest-services-hyper-v-integration-service>

Create and configure generation 1 and generation 2 virtual machines

Windows Server 2012 R2 introduces generation 2 VMs, a special type of VM that differs in configuration from the VMs that are now termed generation 1 VMs, which could be created on Hyper-V virtualization hosts running the Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012 operating systems.

Generation 2 VMs provide the following new functionality:

- Can boot from a SCSI virtual hard disk
- Can boot from a SCSI virtual DVD
- Supports UEFI firmware on the VM
- Supports VM Secure Boot
- PXE boot using standard network adapter

There are no legacy network adapters with generation 2 VMs and the majority of legacy devices, such as COM ports and the Diskette Drive, are no longer present. Generation 2 VMs are “virtual first” and are not designed to simulate hardware for computers that have undergone physical to virtual (P2V) conversion. If you need to deploy a VM that requires an emulated component such as a COM port, you’ll need to deploy a generation 1 VM.

You configure the generation of a VM during VM creation (shown in Figure 1e). Once a VM is created, Hyper-V on Windows Server 2012 R2 doesn’t allow you to modify the VM’s generation. Windows Server 2012 R2 supports running both generation 1 and generation 2 VMs.

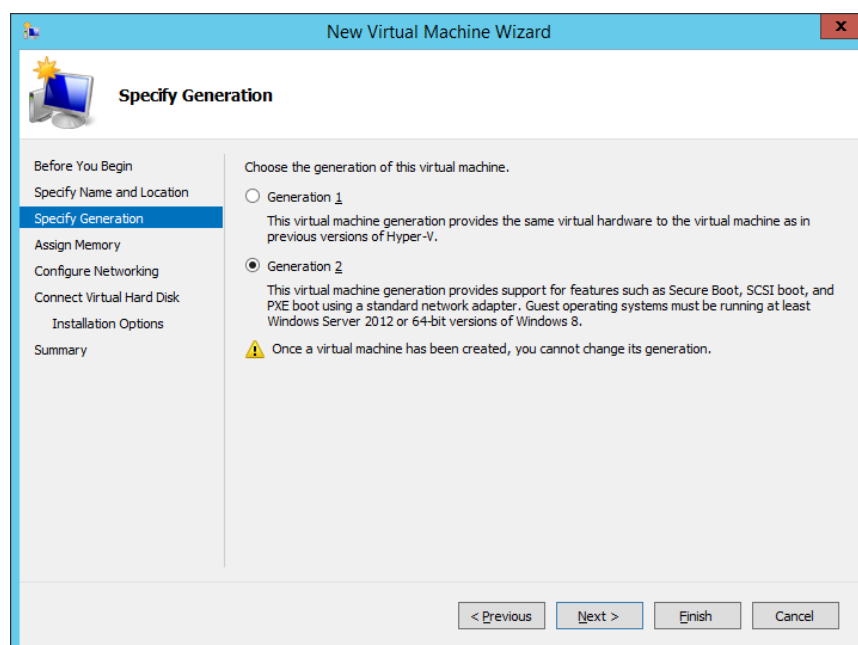


Figure 1e. Select VM generation

Generation 2 VMs boot more quickly and allow the installation of operating systems more quickly than generation 1 VMs. Generation 2 VMs have the following limitations:

- You can only use generation 2 VMs if the guest operating system is running an x64 version of Windows 8, Windows 8.1, Windows Server 2012 or Windows Server 2012 R2
- Generation 2 VMs do not support using RemoteFX
- Generation 2 VMs only support virtual hard disks in VHDX format

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn282285.aspx>

Configure and use Enhanced Session Mode

If you've ever tried to cut and paste text, images or files into a VM when using Virtual Machine Connection when connected to a Hyper-V virtualization host running Windows Server 2008 R2 or Windows Server 2012, you'll know that, except under specific conditions where you can paste text, it's not possible.

Enhanced Session Mode is a feature new to Windows Server 2012 R2 and Windows 8.1 that allows you to perform actions including cutting and pasting, audio redirection, volume and device mapping when using Virtual Machine Connection windows. You can also sign on to a VM with a smart card through Enhanced Session Mode. You enable Enhanced Session Mode on the Hyper-V server by selecting the checkbox for **Allow Enhanced Session Mode** on the Enhanced Session Mode Policy section of the Hyper-V server's properties (shown in Figure 1f).

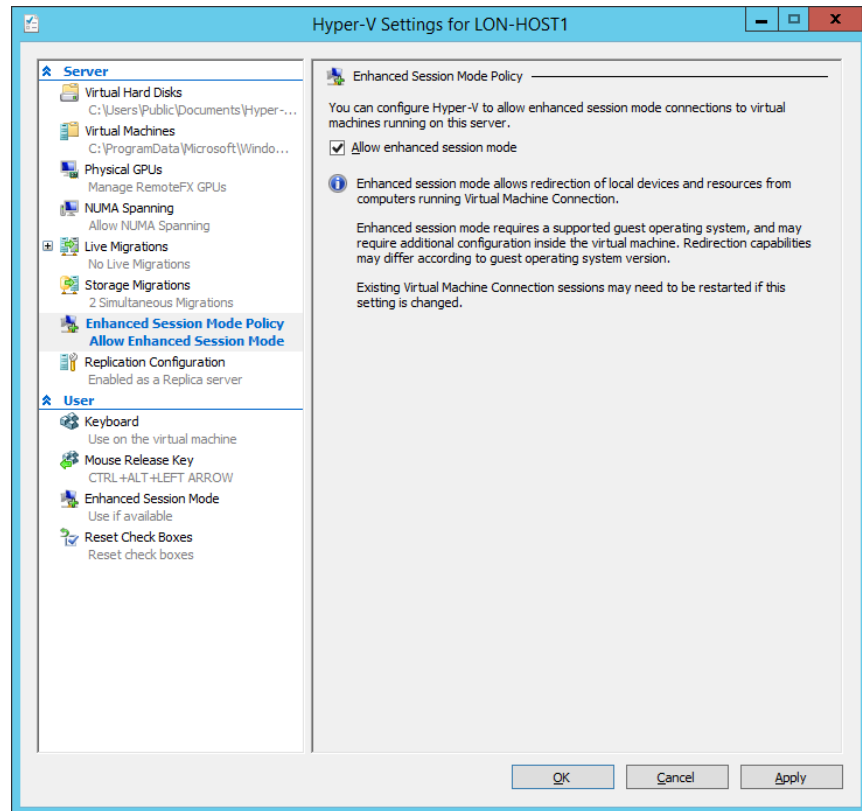


Figure 1f. Select VM generation

Ensure that the checkbox is selected for **Use Enhanced Session Mode** on the Enhanced Session Mode section of the Hyper-V server's properties (shown in Figure 1g).

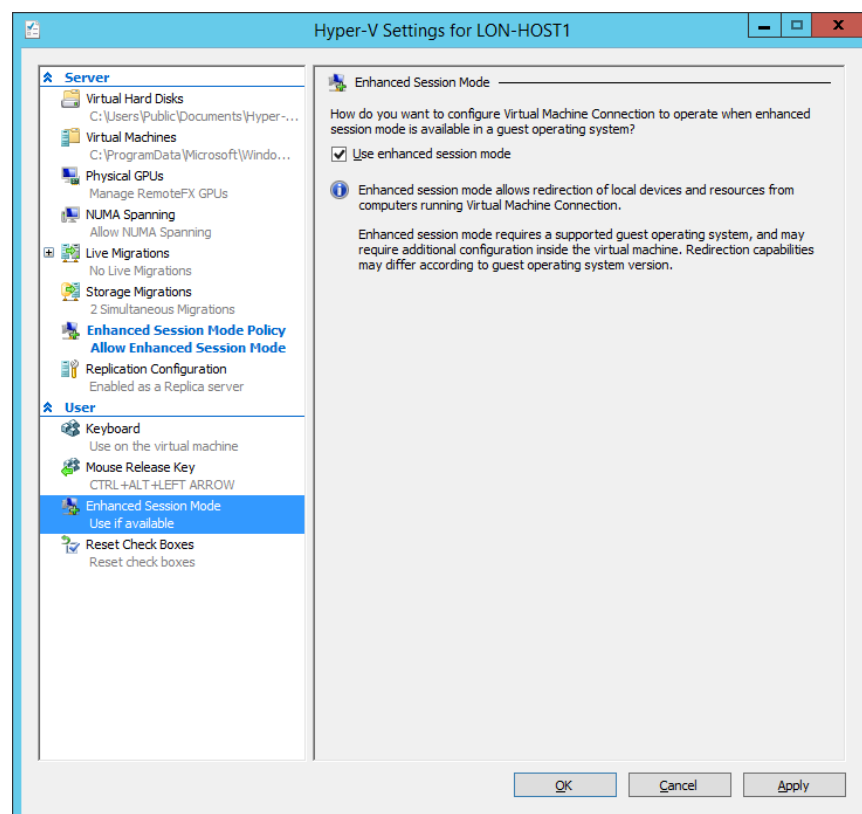


Figure 1g. Select VM generation

You can only use Enhanced Session Mode with guest VMs running the Windows Server 2012 R2 and Windows 8.1 operating systems. To utilize Enhanced Session Mode, you must have permission to connect to the VM using Remote Desktop through the account you use to sign on to the guest VM. You can gain permission by adding the user to the Remote Desktop Users group. A user who is a member of the local Administrators group also has this permission. The Remote Desktop Services service must be running on the guest VM.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn282274.aspx>

Configure RemoteFX

RemoteFX provides a 3D virtual adapter and USB redirection support for VMs. You can only use RemoteFX if the virtualization host has a compatible GPU. RemoteFX allows one or more compatible graphics adapters to perform graphics processing tasks for multiple VMs. You can only use RemoteFX with generation 1 VMs.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/ff817578%28v=ws.10%29.aspx>

Summary

In this chapter, you learned the following:

- Dynamic memory allows you to allocate RAM to VMs based on need rather than configuring a single static amount of RAM. Startup RAM specifies the amount required to start the VM. Minimum RAM is the minimum amount of memory that will be allocated to the VM when there is contention for memory. Maximum RAM is the maximum amount of memory that will be allocated to the VM.
- Integration services allow the exchange of information between the virtualization host and the VM. They also allow the virtualization host to perform operations, such as initiating VM shutdown. Integration services need to be installed on the VM before this functionality is available.
- Smart paging is a technology that simulates extra RAM using files on a disk. It is used when a running VM needs to restart and the configured Startup RAM exceeds the amount of available RAM on the virtualization host.

- Resource metering allows you to track CPU, memory, disk and network utilization on a per-VM basis. You enable resource metering and view resource metering data using Windows PowerShell cmdlets.
- Generation 2 VMs are new to Windows Server 2012 R2 and allow booting of virtual SCSI adapters. Generation 2 VMs use VHDX format virtual hard disks and can only be used with guest operating systems running x64 versions of Windows Server 2012, Windows Server 2012 R2, Windows 8 and Windows 8.1
- Enhanced Session Mode allows you to use Virtual Machine Connection to stream audio from and copy and paste to guest VMs running the Windows 8.1 or Windows Server 2012 R2 operating systems. Enhanced Session Mode also allows you to sign on to a guest VM using a smart card.
- RemoteFX allows you to emulate a remote GPU when compatible hardware is installed on the virtualization host.

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. A VM is configured to use dynamic memory. It is set to use 4 GB of RAM to start and to use a minimum of 3 GB of RAM when it is running. The virtualization host is low on memory and only has 500 MB of RAM available while all VMs are running. The running VM is currently allocated 3 GB of RAM. Which Hyper-V feature would allow this VM to restart after the application of software updates?
2. Describe the changes you can make to the Minimum RAM and Maximum RAM values while a VM is running.
3. Which integration services must be enabled to allow you to copy files from the virtualization host to the VM using the **Copy-VMFile** Windows PowerShell cmdlet.
4. Which Windows PowerShell cmdlet would you use to configure a Hyper-V virtualization host to record CPU, disk, memory and network utilization on a per-VM basis?
5. You want to boot a VM from a virtual hard disk attached to a VM's SCSI controller. What type of VM must you create to accomplish this and what limitations are there on this VM?
6. A user can connect to a VM running Windows 8.1 through Virtual Machine Connection, but is unable to listen to audio. You are able to connect to the same VM from the user's computer and sign on using your own credentials and experience audio playback. What step should you take to resolve this situation?

Chapter 2: Virtual Machine Storage

When you think about storage for your own servers, you probably consider cost, performance, expandability and redundancy before you think about the important specifics of configuring storage. For the 74-409 exam, you'll need to understand considerations around storage for virtualization hosts as well as storage for the virtual machines (VMs) running on those hosts. You need to know how to configure different types of storage and which storage options you would implement given a set of conditions in an exam question scenario.

This chapter covers the following 74-409 exam objectives:

- Create and configure VM storage
- Implement virtualization storage

In this chapter you will learn about:

- Creating virtual hard disks in VHD and VHDX format
- Configuring differencing drives
- Modifying virtual hard disks
- Configuring pass-through disks
- Managing checkpoints
- Implementing virtual Fibre Channel adapters
- Configuring storage Quality of Service
- Configuring Hyper-V host clustered storage
- Configuring guest cluster storage
- Planning for storage optimization

Creating virtual hard disks in VHD and VHDX format

Hyper-V on Windows Server 2012 and Windows Server 2012 R2 supports two separate virtual hard disk formats. The .vhd format has been used with Microsoft virtualization products since the days of Virtual Server. The .vhdx format is new to Windows Server 2012.

Virtual hard disk files in .vhd format are limited to 2040 GB. Virtual hard disks in this format are supported on Hyper-V hosts running Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 and Windows Server 2012 R2. Other than the size limitation, the important thing to remember is that you cannot use virtual hard disk files in .vhd format with generation 2 VMs.

Virtual hard disk files in .vhdx format are an improvement over virtual hard disks in .vhd format. The main limitation of virtual hard disks in .vhdx format is that they cannot be used with Hyper-V on Windows Server 2008 or Windows Server 2008 R2. Virtual hard disks in .vhdx format have the following benefits:

- Can be up to 64 TB in size
- Have larger block size for dynamic and differential disks
- Provide 4-KB logical sector virtual disks
- Have an internal log that reduces chance of corruption
- Support trim to reclaim unused space

You will learn how to convert between virtual hard disks in .vhd and .vhdx format later in this chapter.

You can create virtual hard disks at the time you create the VM, or you can use the New Virtual Hard Disk Wizard (shown in Figure 2a) or the **New-VHD** Windows PowerShell cmdlet.

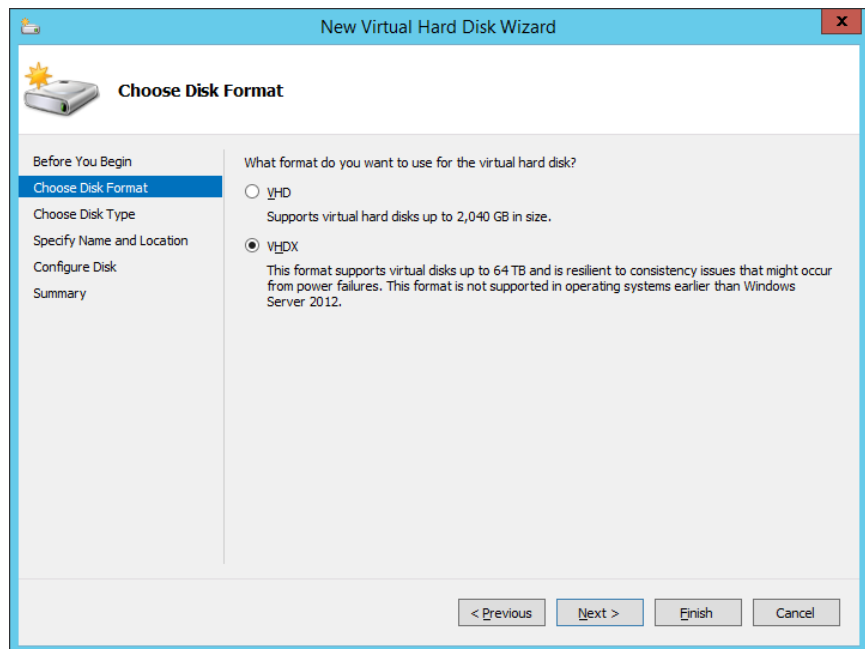


Figure 2a. Select disk format

As shown in Figure 2b, virtual hard disks can either be dynamic, differencing or fixed in size. When you create a fixed-size disk, all space used by the disk is allocated on the hosting volume at the time of creation. Fixed disks increase performance if the physical storage medium does not support Windows Offloaded Data Transfer. Improvements in Windows Server 2012 and Windows Server 2012 R2 reduce the performance benefit of fixed-size disks when the storage medium does support Windows Offloaded Data Transfer. The space to be allocated to the disk must be present on the host volume when you create the disk. For example, you can't create a 3-TB fixed disk on a volume that only has 2 TB of space.

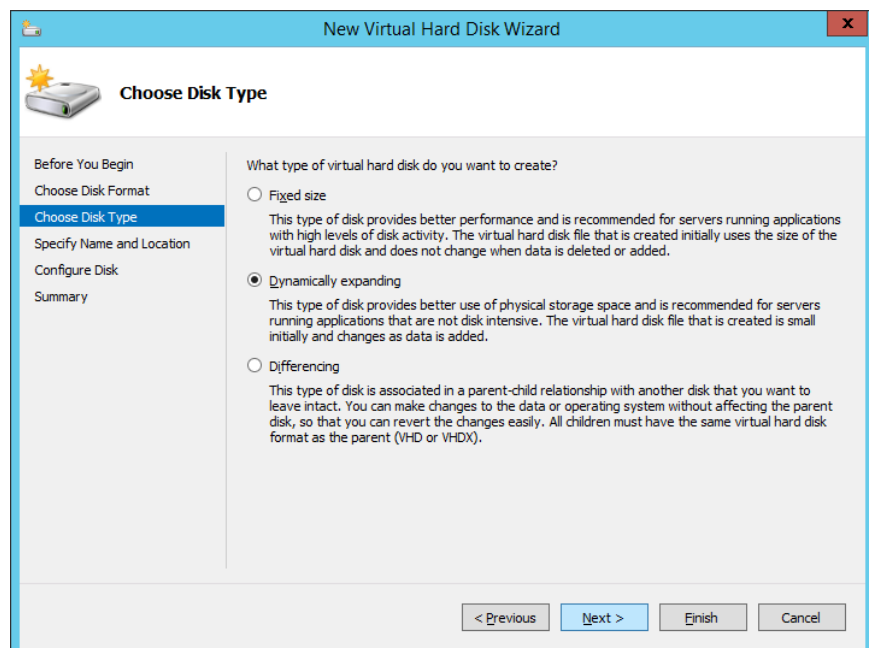


Figure 2b. Select disk type

Dynamically expanding disks use an initial small file and then grow as the VM allocates data to the virtual hard disk. This means you can create a 3-TB dynamic virtual hard disk on a 2-TB volume as the entire 3 TB will not be allocated at disk creation. However, in this scenario you would need to ensure that you extend the size of the 2-TB volume before the dynamic virtual disk outgrows the available storage space.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/hh831446.aspx>

Configuring differencing drives

Differencing disks are a special type of hard disk that is in a child relationship with a parent hard disk. Parent disks can be fixed size or dynamic virtual hard disks, but the differencing disk must be the same type as the parent disk. For example, you can create a differencing disk in .vhdx format for a parent disk that uses .vhdx format, but you cannot create a differencing disk in .vhd format for a parent disk in .vhdx format.

Differencing disks record the changes that would otherwise be made to the parent hard disk by the VM. For example, differencing disks are used to record Hyper-V VM checkpoints. A single parent virtual hard disk can have multiple differencing disks associated with it.

For example, you can create a specially prepared parent virtual hard disk by installing Windows Server 2012 R2 on a VM, running the sysprep utility within the VM and then shutting the VM down. You can use the virtual hard disk created by this process as a parent virtual hard disk. In this scenario, when creating new Windows Server 2012 R2 VMs, you would configure the VMs to use a new differencing disk that uses the sysprepped virtual hard disk as a parent. When you run the new VM, the VM will write any changes that it would make normally to the full virtual hard disk to the differencing disk. In this scenario, deploying new Windows Server 2012 R2 VMs becomes simply a matter of creating new VMs that use a differencing disk that uses the sysprepped Windows Server 2012 R2 virtual hard disk as a parent.

Figure 2c shows the creation of a differencing disk using the New Virtual Hard Disk Wizard where you specify the parent hard disk. You can create differencing hard disks using the **New-VHD** Windows PowerShell cmdlet.




Figure 2c. Creating differencing disk

The key to using differencing disks is to ensure that you don't make changes to the parent disk as this will invalidate the relationship with any child disks. Generally, differencing disks can provide storage efficiencies as the only changes are recorded on child disks. For example, rather than storing 10 different instances of Windows Server 2012 R2 in its entirety, you could create one parent disk and have 10 much smaller differencing disks to accomplish the same objective. If you store VM virtual hard disks on a volume that has been deduplicated, these efficiencies are reduced. You will learn about deduplication later in this chapter.

Modifying virtual hard disks

You can perform the following tasks to modify existing virtual hard disks:

- Convert a virtual hard disk in .vhd format to .vhdx format as shown in Figure 2d
- Convert a virtual hard disk in .vhdx format to .vhd format

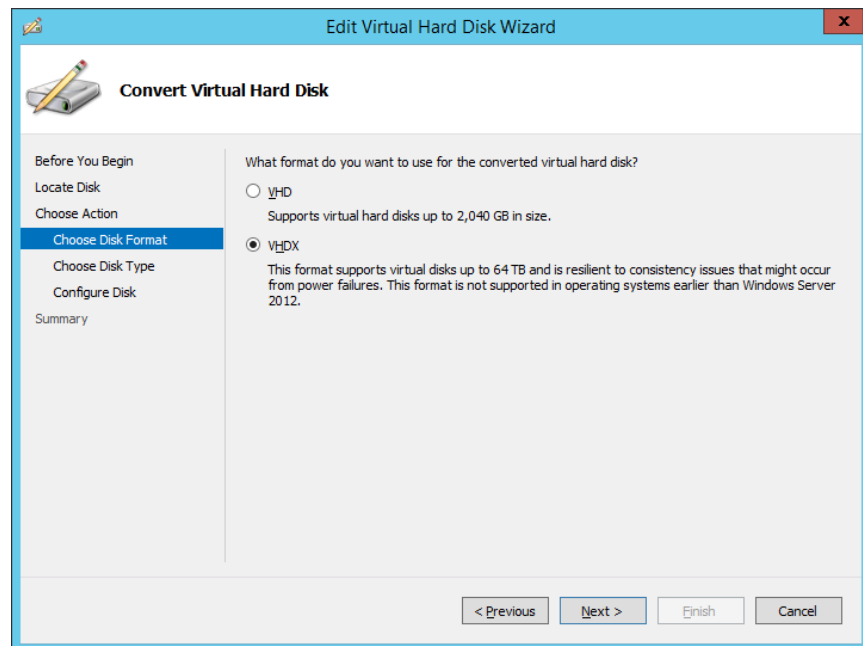


Figure 2d. Convert from VHD to VHDX

- Change the disk from fixed size to dynamically expanding, or from dynamically expanding to fixed size as shown in Figure 2e
- Shrink or enlarge the virtual hard disk

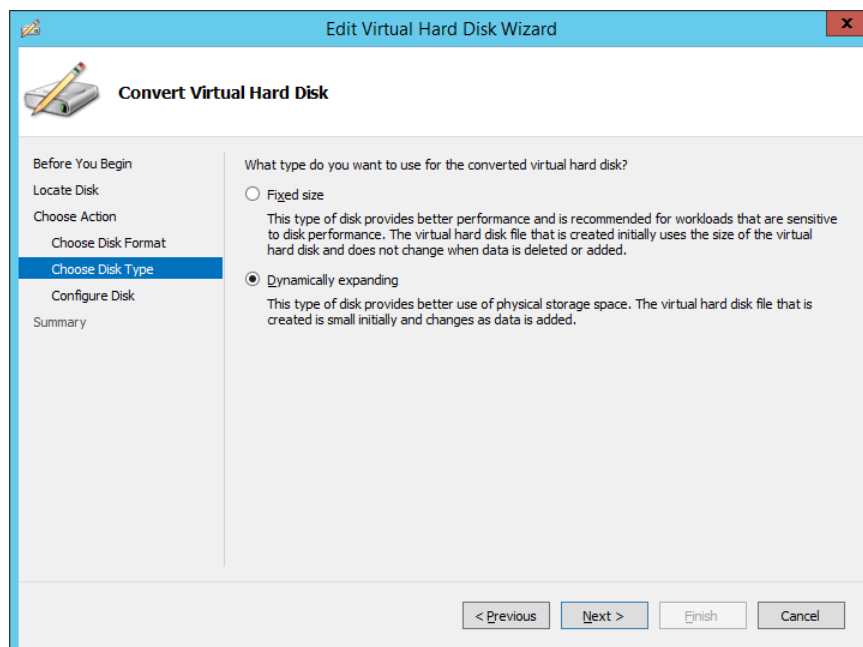


Figure 2e. Convert from fixed to dynamically expanding

You convert virtual hard disk type (.vhd to .vhdx, .vhdx to .vhdx, dynamic to fixed, fixed to dynamic) either using the Edit Virtual Hard Disk Wizard or by using the **Convert-VHD** Windows PowerShell cmdlet. When converting from .vhdx to .vhd, remember that virtual hard disks in .vhd format cannot exceed 2040 GB in size. So while it is possible to convert virtual hard disks in .vhdx format that are smaller than 2040 GB to .vhd format, you will not be able to convert virtual hard disks that are larger than this.

You can only perform conversions from one format to another and from one type to another while the VM is powered off. You must shrink the virtual hard disk using disk manager in the VM operating system prior to shrinking the virtual hard disk using the Edit Virtual Hard Disk Wizard or **Resize-VHD** cmdlet. You can resize a virtual hard disk while the VM is running under the following conditions:

- The virtualization host is running Windows Server 2012 R2
- The virtual hard disk is in .vhdx format
- The virtual hard disk is attached to a virtual SCSI controller
- The virtual hard disk must have been shrunk
 - You must shrink the virtual hard disk using disk manager in the host operating system prior to shrinking the virtual hard disk using the Edit Virtual Hard Disk Wizard or **Resize-VHD** cmdlet

Find Out More:

You can learn more about performing online resizing virtual hard disks by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn282286.aspx>

Configuring pass-through disks

Pass-through disks, also known as directly attached disks, allow a VM direct access to underlying storage rather than a virtual hard disk that resides on that storage. For example, normally with Hyper-V you will connect a VM to a virtual hard disk file hosted on a volume formatted with NTFS. With pass-through disks, the VM instead accesses the disk directly and there is no virtual hard disk file.

Pass-through disks allow VMs to access larger volumes than are possible when using virtual hard disks in .vhd format. In earlier versions of Hyper-V, such as the version available with Windows Server 2008, pass-through disks provided performance advantages over virtual hard disks. The need for pass-through disks has diminished with the availability of virtual hard disks in .vhdx format, as .vhdx format allows you to create much larger volumes.

Pass-through disks can be directly attached to the virtualization host or can be Fibre Channel or iSCSI disks as shown in Figure 2f. When adding a pass-through disk, you will need to ensure that the disk is offline. You can use the disk management console or the diskpart.exe utility on the virtualization host to set a disk to be offline.

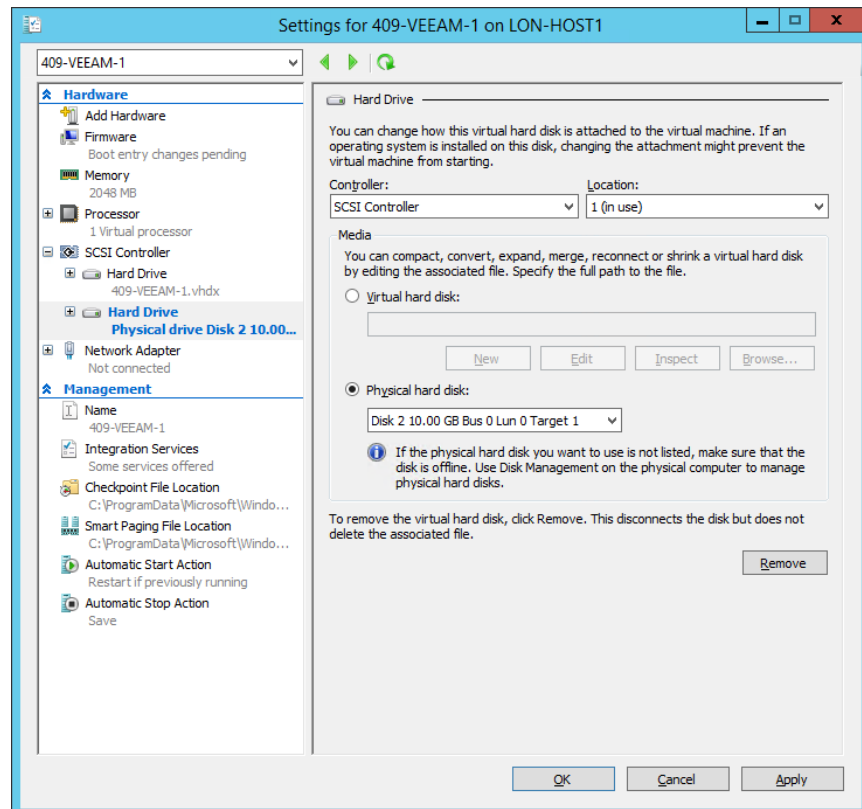


Figure 2f. Connecting pass-through disk

To add a pass-through disk using Windows PowerShell, first get the properties of the disk that you want to add as a pass-through disk using the `Get-Disk` cmdlet and then pipe the result to the `Add-VMHardDiskDrive` cmdlet. For example, to add physical disk 3 to the VM named `Veeam-Test`, execute the following command:

```
Get-Disk 3 | Add-VMHardDiskDrive -VMName Veeam-Test
```

A VM that uses pass-through disks will not support VM checkpoints. Pass-through disks also cannot be backed up using any backup program that uses the Hyper-V VSS writer.

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/askcore/archive/2008/10/24/configuring-pass-through-disks-in-hyper-v.aspx>. This article deals with an earlier version of the technology, but the information it contains remains relevant to pass-through disks in Windows Server 2012 R2.

Managing checkpoints

Checkpoints represent the state of a VM at a particular point in time. In previous versions of Hyper-V, checkpoints were known as snapshots. You can create checkpoints when the VM is running or when the VM is shut down. When you create a checkpoint of a running VM, the running VM's memory state is also stored in the checkpoint. Restoring a checkpoint taken of a running VM returns the running VM to a restored state. Creating a checkpoint creates either an .avhd or .avhdx file (depending on whether the VM is using virtual hard disks in the VHD or VHDX format).

You can create checkpoints from Windows PowerShell with the **Checkpoint-VM** cmdlet. The other checkpoint-related Windows PowerShell cmdlets actually use the **VMSnapshot** noun. These cmdlets are as follows:

- **Restore-VMSnapshot.** Restores an existing VM checkpoint.
- **Export-VMSnapshot.** Allows you to export the state of a VM as it exists when a particular checkpoint was taken. For example, if you took a checkpoint at 2 p.m. and 3 p.m., you could choose to export the checkpoint taken at 2 p.m. and then import the VM in the state that it was in at 2 p.m. on another Hyper-V host.
- **Get-VMSnapshot.** Lists the current checkpoints.
- **Rename-VMSnapshot.** Allows you to rename an existing VM checkpoint.
- **Remove-VMSnapshot.** Deletes a VM checkpoint. If the VM checkpoint is part of the chain, but not the final link, changes are merged with the successive checkpoint so that it remains a representation of the VM at the point in time when the snapshot was taken. For example, if checkpoints were taken at 1 p.m., 2 p.m. and 3 p.m. and you delete the 2 p.m. checkpoint, the avhd/avhdx files associated with the 2 p.m. snapshot would be merged with the avhd/avhdx files associated with the 3 p.m. snapshot so that the 3 p.m. snapshot retained its integrity.

Checkpoints do not replace backups. Checkpoints are almost always stored on the same volume as the original VM hard disks, so a failure of that volume will result in all VM storage files, both original disks and checkpoint disks, being lost. If a disk in a checkpoint chain becomes corrupted, then that checkpoint and all subsequent checkpoints will be lost. Disks earlier in the checkpoint chain will remain unaffected.

Hyper-V supports a maximum of 50 checkpoints per VM. When using System Center 2012 R2 Virtual Machine Manager, up to 64 checkpoints per VM are supported.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dd851843.aspx>. The article deals with Windows Server 2008 R2 rather than Windows Server 2012 and Windows Server 2012 R2, but provides solid background information on snapshots, which evolved into checkpoints.

Implementing virtual Fibre Channel adapters

Virtual Fibre Channel allows you to make direct connections from VMs running on Hyper-V to Fibre Channel storage. Virtual Fibre Channel is supported on Windows Server 2012 and Windows Server 2012 R2 if the following requirements are met:

- The computer functioning as the Hyper-V virtualization host must have a Fibre Channel host bus adapter (HBA) that has a driver that supports virtual Fibre Channel.
- SAN must be NPIV (N_Port ID) enabled.
- The VM must be running Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2 as the guest operating system.
- Virtual Fibre Channel logical units cannot be used to boot Hyper-V VMs

VMs running on Hyper-V support up to four virtual Fibre Channel adapters, each of which can be associated with a separate Storage Area Network (SAN).

Before you can use a virtual Fibre Channel adapter, you will need to create at least one virtual SAN on the Hyper-V virtualization host. A virtual SAN is a group of physical Fibre Channel ports that connect to the same SAN.

VM live migration and VM failover clusters are supported; however, virtual Fibre Channel does not support VM checkpoints, host-based backup or live migration of SAN data.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831413.aspx>

Configuring Storage Quality of Service

Storage Quality of Service (QoS) allows you to limit the maximum number of IOPS (Input/Output operations per second) for virtual hard disks. IOPS are measured in 8-KB increments. If you specify a maximum IOPS value, the virtual hard disk will be unable to exceed this value. You use Storage QoS to ensure that no one single workload on a Hyper-V virtualization host consumes a disproportionate amount of storage resources.

It's also possible to specify a minimum IOPS value for each virtual hard disk. You would do this if you wanted to be notified that a specific virtual hard disk's IOPS has fallen below a threshold value. When the number of IOPS falls below the specified minimum, an event is written to the event log. You configure Storage QoS on a per-virtual hard disk basis as shown in Figure 2g.

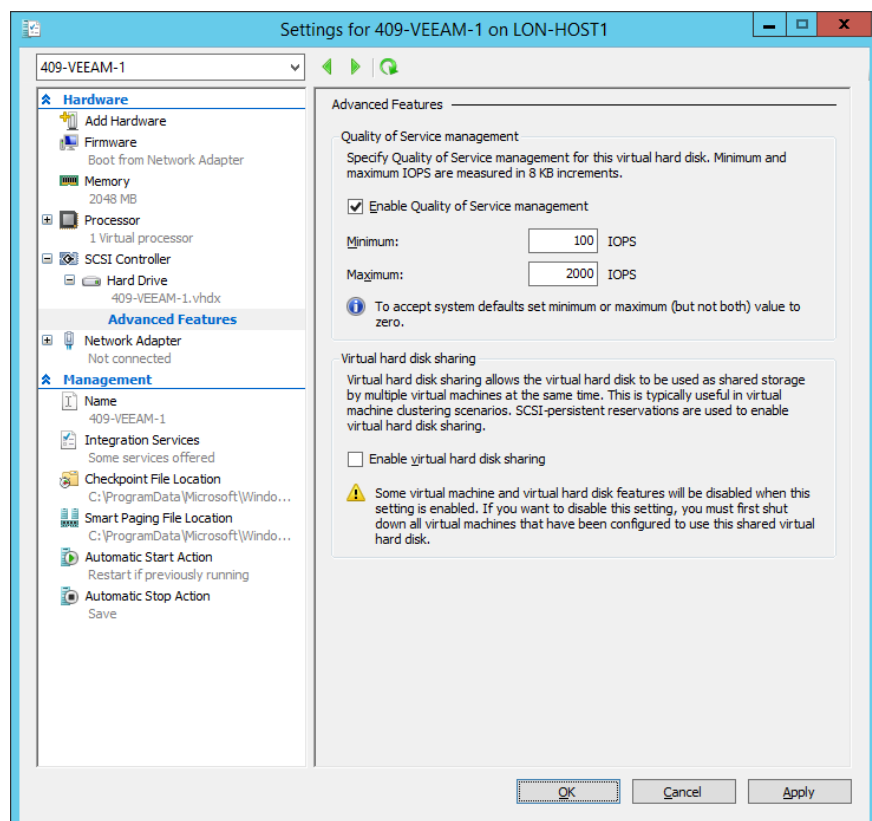


Figure 2g. Configuring QoS

You cannot configure Storage QoS if you have enabled a virtual hard disk for virtual hard disk sharing.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/library/dn282281.aspx>

Configuring Hyper-V host clustered storage

When deployed on Hyper-V host clusters, the configuration and virtual hard disk files for highly available VMs are hosted shared storage. This shared storage can be one of the following:

- **Serial Attached SCSI (SAS).** Suitable for two-node failover clusters where the cluster nodes are in close proximity to each other.
- **iSCSI storage.** Suitable for failover clusters with two or more nodes. Windows Server 2012 and Windows Server 2012 R2 includes iSCSI Target Software, allowing it to host iSCSI targets that can be used as shared storage by Windows failover clusters.
- **Fibre Channel.** Fibre Channel/Fibre Channel over Ethernet storage requires special network hardware. While generally providing better performance than iSCSI, Fibre Channel components tend to be more expensive.
- **SMB 3.0 file shares configured as continuously available storage.** This special type of file share is highly available, with multiple cluster nodes able to maintain access to the file share. This configuration requires multiple clusters. One cluster hosts the highly available storage used by the VMs, and the other cluster hosts the highly available VMs. Scale out file servers using SMB 3.0 are only available using Windows Server 2012 and Windows Server 2012 R2.
- **Cluster Shared Volumes (CSVs).** CSVs can also be used for VM storage in Hyper-V failover clusters. As with continuously available file shares, multiple nodes in the cluster have access to the files stored on CSVs, ensuring that failover occurs with minimal disruption. As with SMB 3.0 file shares, multiple clusters are required, with one cluster hosting the CSVs and the other cluster hosting the VMs. CSVs can be used as shared storage for Windows Server 2008 R2 Hyper-V failover clusters.

When considering storage for a Hyper-V failover cluster, remember the following:

- Keep in mind that Windows Server 2012 and Windows Server 2012 R2 do not support dynamic disks as shared storage.
- Ensure volumes used for disk witnesses are formatted as either NTFS or ReFS.
- Avoid allowing nodes from separate failover clusters to access the same shared storage by using LUN masking or zoning.
- Where possible, use storage spaces to host volumes presented as shared storage.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/library/jj612869.aspx>

Configuring guest cluster storage

Just as you can configure a Hyper-V failover cluster where multiple Hyper-V hosts function as failover cluster nodes, you can configure failover clusters within VMs, where each failover cluster node is a VM. Even though failover cluster nodes must be members of the same Active Directory domain, there is no requirement that they be hosted on the same cluster. For example, you could configure a multi-site failover cluster where the cluster nodes are hosted as highly available VMs, each hosted on its own Hyper-V failover clusters in each site.

When considering how to deploy a VM guest cluster, you will need to choose how you will provision the shared storage that is accessible to each cluster node. The options for configuring shared storage for VM guest clusters include:

- iSCSI
- Virtual Fibre Channel
- Cluster Shared Volumes
- Continuously Available File Shares
- Shared virtual hard disks

The conditions for using iSCSI, Virtual Fibre Channel, Cluster Shared Volumes and Continuously Available File Shares with VM guest clusters are essentially the same for VMs as they are when configuring traditional physically hosted failover cluster nodes.

Shared virtual hard disks are a special type of shared storage only available to VM guest clusters. With shared virtual hard disks, each guest cluster node can be configured to access the same shared virtual hard disk. Each VM cluster node's operating system will recognize the shared virtual hard disk as shared storage when building the VM guest failover cluster. Figure 2h shows a virtual hard disk being configured as a shared virtual hard disk.

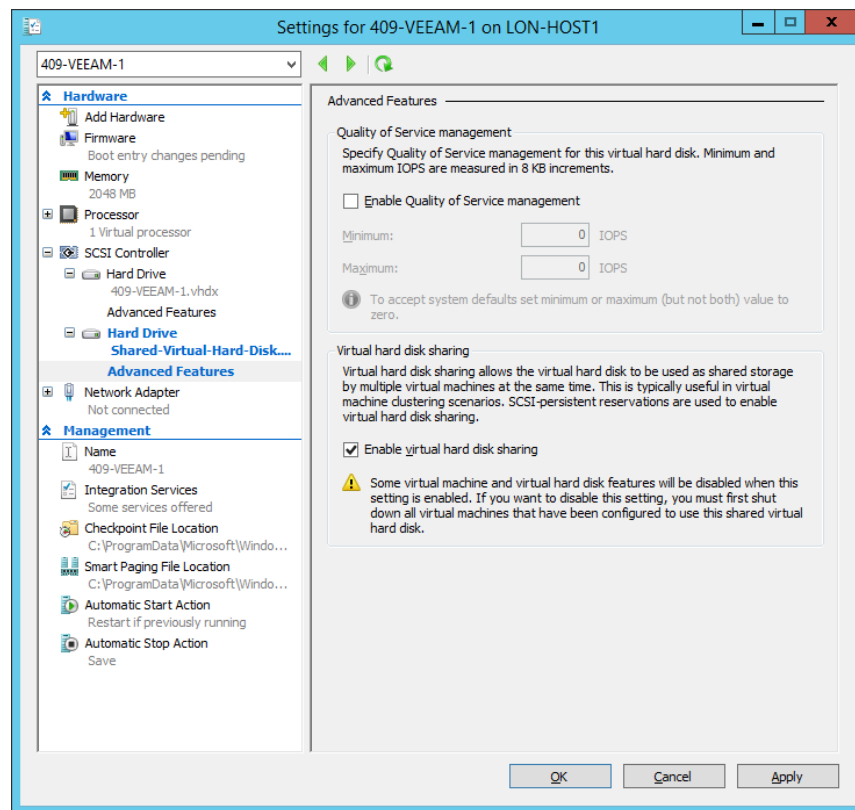


Figure 2h. Configure shared virtual hard disk

Shared virtual hard disks have the following requirements:

- Can be used with generation 1 and generation 2 VMs.
- Can only be used with guest operating systems running Windows Server 2012 or Windows Server 2012 R2. If the guest operating systems are running Windows Server 2012, they must be updated to use the Windows Server 2012 R2 integration services components.
- Can only be used if virtualization hosts are running the Windows Server 2012 R2 version of Hyper-V.
- Must be configured to use the .vhdx virtual hard disk format.
- Must be connected to a virtual SCSI controller.
- When deployed on a failover cluster, the shared virtual hard disk itself should be located on shared storage, such as a Continuously Available File Share or Cluster Shared Volume. This is not necessary when configuring a guest failover cluster on a single Hyper-V server that is not part of a Hyper-V failover cluster.
- VMs can only use shared virtual hard disks to store data. You can't boot a VM from a shared virtual hard disk.

The configuration of shared virtual hard disks differs from the traditional configuration of VM guest failover clusters, because you configure the connection to shared storage by editing the VM properties rather than connecting to the shared storage from within the VM.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/library/dn265980.aspx>

Planning for storage optimization

New technologies built into Windows Server 2012 and Windows Server 2012 R2 allow you to optimize the performance and data storage requirements for files associated with VMs.

Deduplication

NTFS and ReFS volumes on computers running Windows Server 2012 and Windows Server 2012 R2 support deduplication. Deduplication is a process by which duplicate instances of data are removed from a volume and replaced with pointers to the original instance. Deduplication is especially effective when used with volumes that host virtual hard disk files as many of these files contain duplicate copies of data, such as the VM's operating system and program files. Deduplication can be installed as a role feature on computers running the Windows Server 2012 and Windows Server 2012 R2 operating systems as shown in Figure 2i.

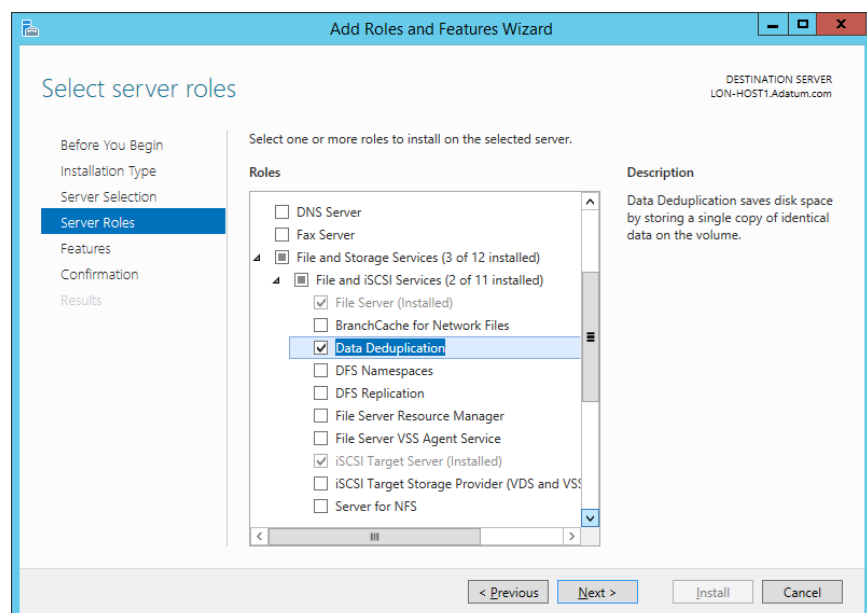


Figure 2i. Data deduplication feature

Once installed, you can enable deduplication through the Volumes node of the File and Storage Services section of the Server Manager Console. When enabling deduplication, you specify whether you want to use a general file server data deduplication scheme or a virtual desktop infrastructure scheme. For volumes that host VM files, the VDI scheme is appropriate. Figure 2j shows enabling deduplication using the VDI scheme.

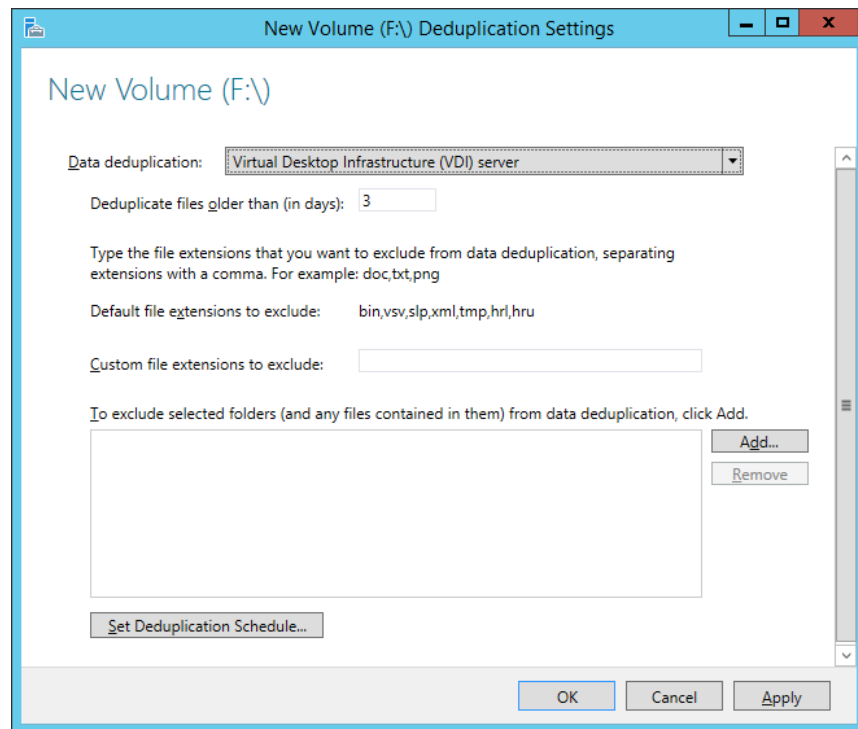


Figure 2j. Enabling deduplication

You can't enable deduplication on the operating system volume, only on data volumes.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831602.aspx>

Storage tiering

Storage tiering is a technology that allows you to mix fast storage, such as solid state disk (SSD), with traditional spinning magnetic disks to optimize both storage performance and capacity. Storage tiering works on the premise that a minority of the data stored on a volume is responsible for the majority of read and write operations. Rather than creating a large volume that consists entirely of SSDs, storage tiering (through Windows Server 2012 R2's storage spaces functionality) allows you to create a volume comprised of both solid state and spinning magnetic disks. In this configuration, frequently accessed data is moved to the parts of the volume hosted on the SSDs and less frequently accessed data is moved to the parts of the volume hosted on the slower spinning magnetic disks. This configuration allows much of the performance benefits of an SSD-only volume to be realized without the cost of using SSD-only storage.

When used in conjunction with deduplication, frequently accessed deduplicated data is moved to the faster storage, allowing both the benefit of reducing storage requirements while improving performance over what would be possible if the volume hosting VM files were comprised only of spinning magnetic disks. You also have the option of pinning specific files to the faster storage, which overrides the algorithms that move data according to accumulated utilization statistics.

You configure storage tiering using Windows PowerShell.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
http://technet.microsoft.com/en-us/library/dn387076.aspx#bkmk_tiers

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. Which virtual hard disk format would you use if you wanted to provision a VM with 3072-GB boot volume?
2. Under what conditions can you reduce the size of a VM's virtual hard disk while the VM remains in operation?
3. Which Windows PowerShell cmdlet do you use to revert a VM running on Windows Server 2012 R2 Hyper-V to a previously created checkpoint?
4. Which Windows PowerShell cmdlet do you use to create a differencing virtual hard disk?
5. What condition must the virtualization host running Windows Server 2012 R2 meet if you are going to be able to implement virtual Fibre Channel on the VMs it hosts?
6. What functionality would you configure to ensure that no one particular virtual hard disk's IOPS overwhelm a Hyper-V server's storage?
7. In what state must a disk be if you are going to connect it as pass-through storage to a Hyper-V VM?
8. Which operating systems support enabling the deduplication role service on volumes hosting VMs?
9. What conditions can be met before a VM guest cluster running on a Windows Server 2012 R2 virtualization host can use shared virtual hard disk as shared storage?
10. You have 4 TB available on a storage spaces volume. You will be adding an additional 5 TB to this volume in the next few weeks, but you need to create a 6-TB volume for use with a specific VM. This VM will not need all this storage immediately and will only require 2 TB of data to be stored in the next few months. What type of virtual hard disk will you create to accomplish this goal?

Chapter 3: Hyper-V Virtual Networks and Virtualization Networking

Once you've configured virtual machine (VM) settings and storage, the next step is to configure VM networking. You configure networking for a VM by configuring virtual network adapters. You configure networking for an individual Hyper-V server by configuring virtual switches. When you need to configure networking for a large number of Hyper-V servers, you configure VMM logical networks, port profiles, logical switches and VM networks.

In this chapter you'll learn about:

- Hyper-V virtual switches
- Optimizing network performance
- Configuring MAC addresses
- Configure network isolation
- Configure synthetic and legacy virtual network adapters
- VM NIC teaming
- VMM logical networks
- VMM port profiles and logical switches
- VMM VM networks
- VMM IP and MAC address pools
- Windows Server Gateway
- Private Virtual Local Area Networks

Hyper-V virtual switches

Hyper-V virtual switches, termed Hyper-V virtual networks in previous versions of Hyper-V, represent network connections to which the Hyper-V virtual network adapters can connect. You can configure three types of Hyper-V virtual switches:

- **External.** An external switch connects to a physical or wireless network adapter. Only one virtual switch can be mapped to a specific physical or wireless network adapter or NIC team. For example, if a virtualization host had four physical network adapters configured as two separate NIC teams, you could configure two external virtual switches. If a virtualization host had three physical network adapters that did not participate in any NIC teams, you could configure three external virtual switches. Any VMs connected to the same external switch can communicate with each other as well as any external host connected to the network that the network adapter mapped to the external switch is connected to. For example, if an external switch is connected to a network adapter that is connected to a network that can route traffic to the internet, a VM connected to that external virtual switch will also be able to connect to hosts on the internet. Figure 3a shows an external switch mapped to a Realtek PCIe GBE network adapter. When you create an external switch, a virtual network adapter that maps to this switch is created on the virtualization host unless you clear the option that allows the management operating system to share the network adapter. If you clear this option, the virtualization host will not be able to communicate through the network adapter associated with the external switch.

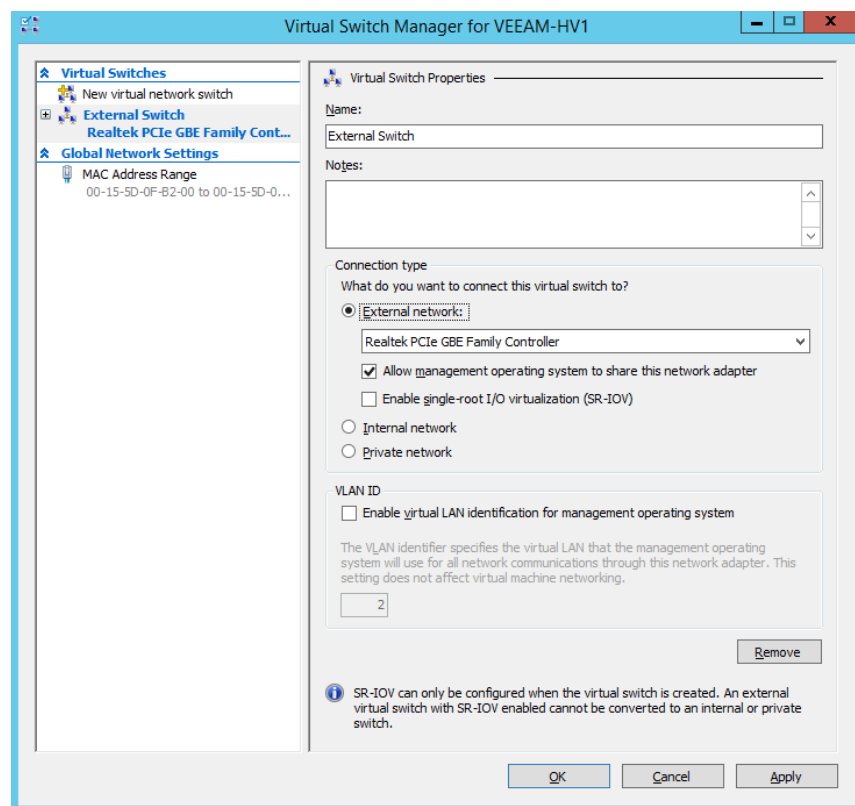


Figure 3a. External virtual switch

- **Internal.** An internal switch allows communication between the VM and the virtualization host. All VMs connected to the same internal switch are able to communicate with each other and the virtualization host. For example, you could successfully initiate an RDP connection from the virtualization host to an appropriately configured VM or use the **Test-NetConnection** Windows PowerShell cmdlet from a Windows PowerShell prompt on the virtualization host to get a response from a VM connected to an internal network connection. VMs connected to an internal switch are unable to use that virtual switch to communicate with hosts on a separate virtualization host that are connected to an internal switch with the same name. Figure 3b shows an internal switch.

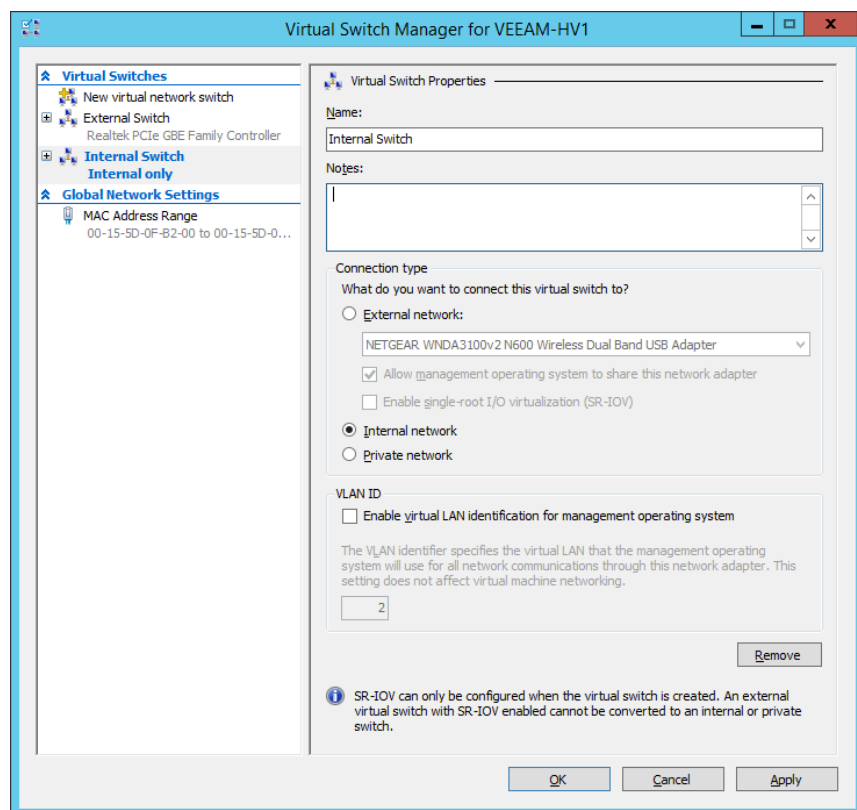


Figure 3b. Internal virtual switch

- **Private.** VMs connected to the same private switch on a VM host are able to communicate with one another, but are unable to communicate directly with the virtualization host. Private switches only allow communication between VMs on the same virtualization host. For example, VM alpha and beta are connected to private switch p_switch_a on virtualization host h_v_one. VM gamma is connected to private switch p_switch_a on virtualization host h_v_two. VMs alpha and beta will be able to communicate with each other, but will be unable to communicate with h_v_one or VM gamma.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/hh831823.aspx>

Optimizing network performance

You can optimize network performance for VMs hosted on Hyper-V in a number of ways. For example, you can configure the virtualization host with separate network adapters connected to separate subnets. You do this to separate network traffic related to the management of the Hyper-V virtualization host from network traffic associated with hosted VMs. You can also use NIC teaming on the Hyper-V virtualization host to provide increased and fault-tolerant network connectivity. You'll learn more about NIC teaming later in the chapter.

An additional method of optimizing network performance is to configure bandwidth management at the virtual network adapter level. Bandwidth management allows you to specify a minimum and a maximum traffic throughput figure for a virtual network adapter. The minimum bandwidth allocation is an amount that Hyper-V will reserve for the network adapter. For example, if you set the minimum bandwidth allocation to 10 Mbps for each VM, Hyper-V would ensure that when other VMs needed more, they would be able to increase their bandwidth utilization until they reached a limit defined by the combined minimum bandwidth allocation of all VMs hosted on the server. Maximum bandwidth allocations specify an upper limit for bandwidth utilization. By default, no minimum or maximum limits are set on virtual network adapters.

You configure bandwidth management by selecting the **Enable bandwidth management** option on a virtual network adapter and specifying a minimum and maximum bandwidth allocation in megabits per seconds (Mbps). Figure 3c shows bandwidth management enabled for a virtual network adapter connected to an external switch, with the minimum bandwidth set to 0 Mbps and the maximum bandwidth set to 100 Mbps.

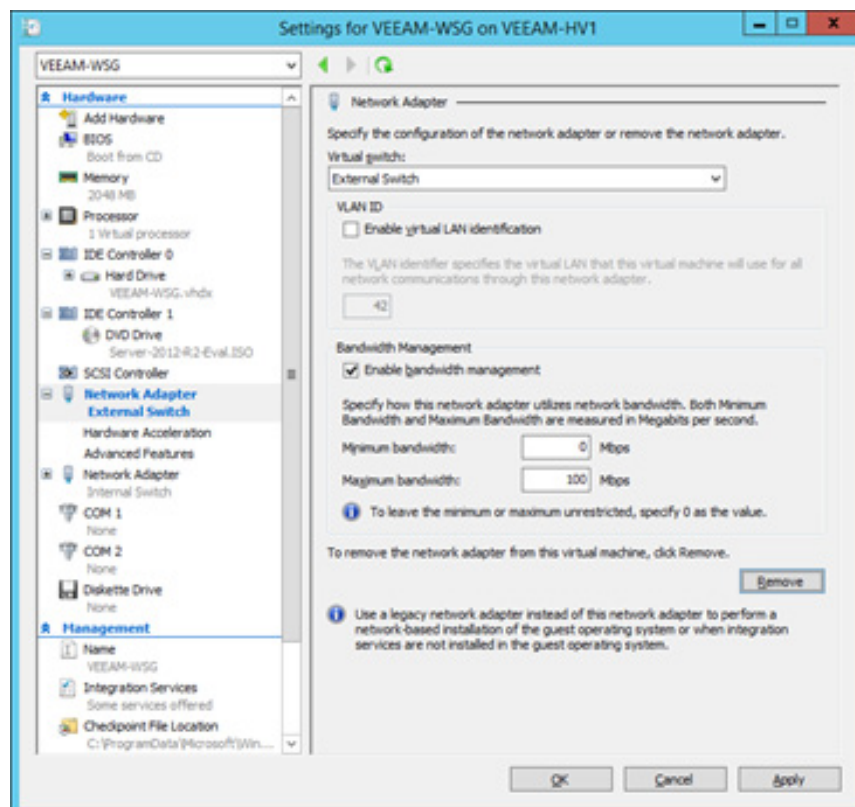


Figure 3c. Bandwidth management

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj735302.aspx>

SR-IOV (Single Root I/O Virtualization) increases network throughput by bypassing a virtual switch and sending network traffic straight to the VM. When you configure SR-IOV, the physical network adapter is mapped directly to the VM. As such, SR-IOV requires that the VM's operating system include a driver for the physical network adapter. You can only use SR-IOV if the physical network adapter and the network adapter drivers used with the virtualization host support the functionality. You can only configure SR-IOV for a virtual switch during switch creation. Once you have an SR-IOV-enabled virtual switch, you can then enable SR-IOV on the virtual network adapter that connects to that switch, as shown in Figure 3d.

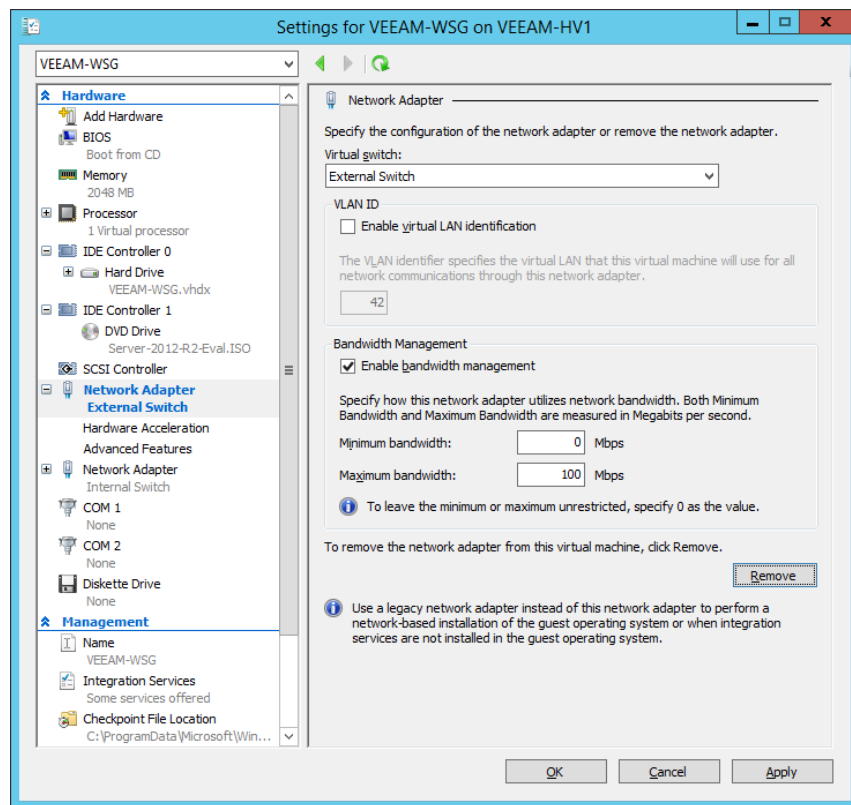


Figure 3d. SR-IOV and virtual machine queue

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/jhoward/archive/2012/03/12/everything-you-wanted-to-know-about-sr-iov-in-hyper-v-part-1.aspx>

Dynamic virtual machine queue is an additional technology that you can use to optimize network performance. When a VM is connected through a virtual switch to a network adapter that supports virtual machine queue and virtual machine queue is enabled on the virtual network adapter's properties, the physical network adapter is able to use Direct Memory Access (DMA) to forward traffic directly to the VM. With virtual machine queue, network traffic is processed by the CPU assigned to the VM rather than by the physical network adapter used by the Hyper-V virtualization host. Dynamic virtual machine queue automatically adjusts the number of CPU cores used to process network traffic. Dynamic virtual machine queue is automatically enabled on a virtual switch when you enable virtual machine queue on the virtual network adapter.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/gg162704%28v=ws.10%29.aspx>

Virtual machine MAC addresses

By default, VMs running on Hyper-V hosts use dynamic MAC addresses. Each time a VM is powered on, it will be assigned a MAC address from a MAC address pool. You can configure the properties of the MAC address pool through the MAC Address Range settings available through Virtual Switch Manager, as shown in Figure 3e.

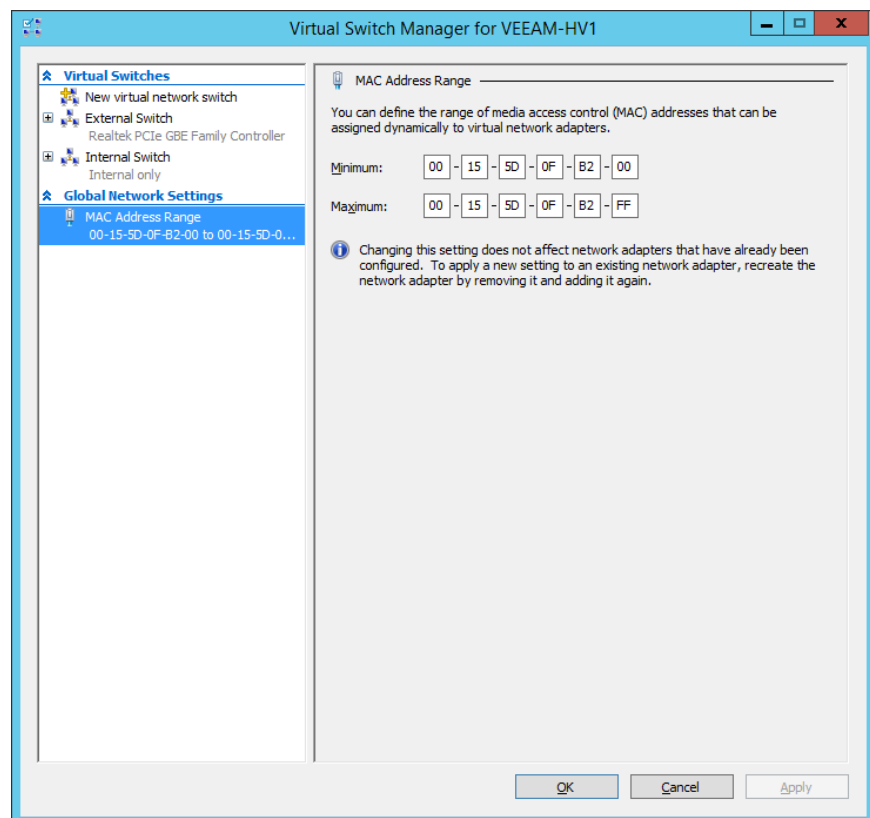


Figure 3e. MAC address pool

When you deploy operating systems on physical hardware, you can use two methods to ensure that the computer is always assigned the same IP address configuration. The first method is to assign a static IP address from within the virtualized operating system. The second is to configure a DHCP reservation that always assigns the same IP address configuration to the MAC address associated with the physical computer's network adapter.

This won't work with Hyper-V VMs in their default configuration because the MAC address may change if you power the VM off and then on. Rather than configure a static IP address using the VM's operating system, you can instead configure a static MAC address on a per-virtual network adapter basis. This will ensure that a VM's virtual network adapter retains the same MAC address whether the VM is restarted or even if the VM is migrated to another virtualization host.

To configure a static MAC address on a per-network adapter basis, edit the network adapter's advanced features, as shown in Figure 3f. When entering a static MAC address, you will need to select a MAC address manually. You shouldn't use one from the existing MAC address pool as there is no way for the current virtualization hosts, or other virtualization hosts on the same subnet, to check whether a MAC address that is to be assigned dynamically has already been assigned statically.

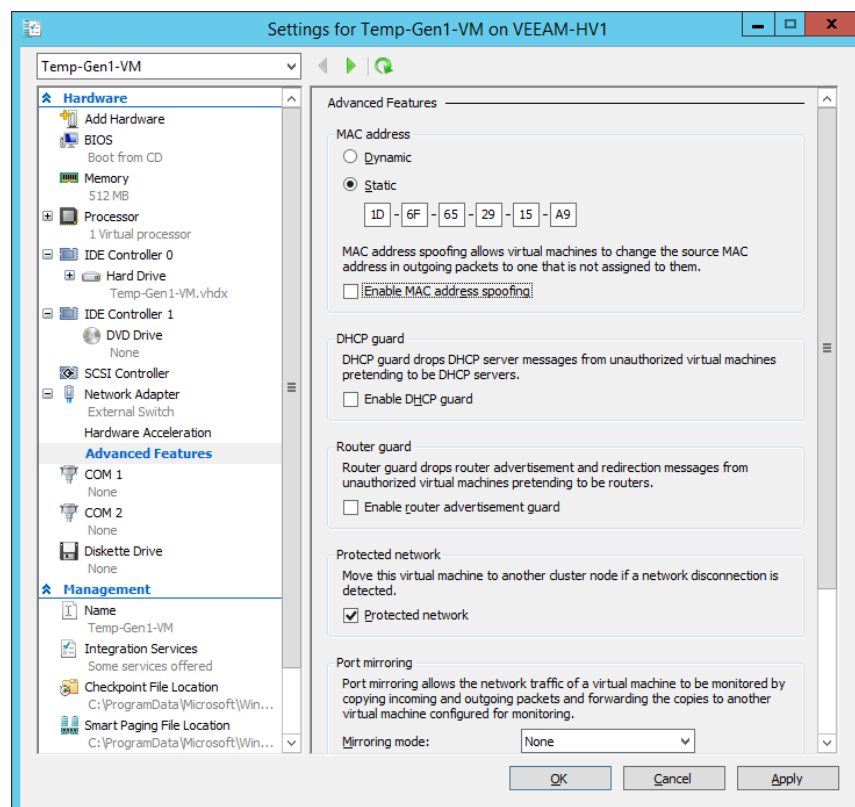


Figure 3f. Configure static MAC address

Find Out More:

You can learn more about Hyper-V MAC addresses by consulting the following TechNet article: <http://blogs.msdn.com/b/robertvi/archive/2011/03/25/how-does-a-nic-inside-a-vm-get-a-mac-address-assigned.aspx?Redirected=true>

Configuring network isolation

Hyper-V supports VLAN (Virtual Local Area Network) tagging at both the network adapter and virtual switch level. VLAN tags allow the isolation of traffic for hosts connected to the same network by creating separate broadcast domains. Enterprise hardware switches also support VLANs as a way of partitioning network traffic. To use VLANs with Hyper-V, the virtualization hosts' network adapter must support VLANs. A VLAN ID has 12 bits, which means you can configure 4,094 VLAN IDs.

You configure VLAN tags at the virtual network adapter level by selecting the **Enable virtual LAN identification** checkbox in the virtual network adapter properties. Figure 3g shows the VLAN identifier for a virtual network adapter set to 42.

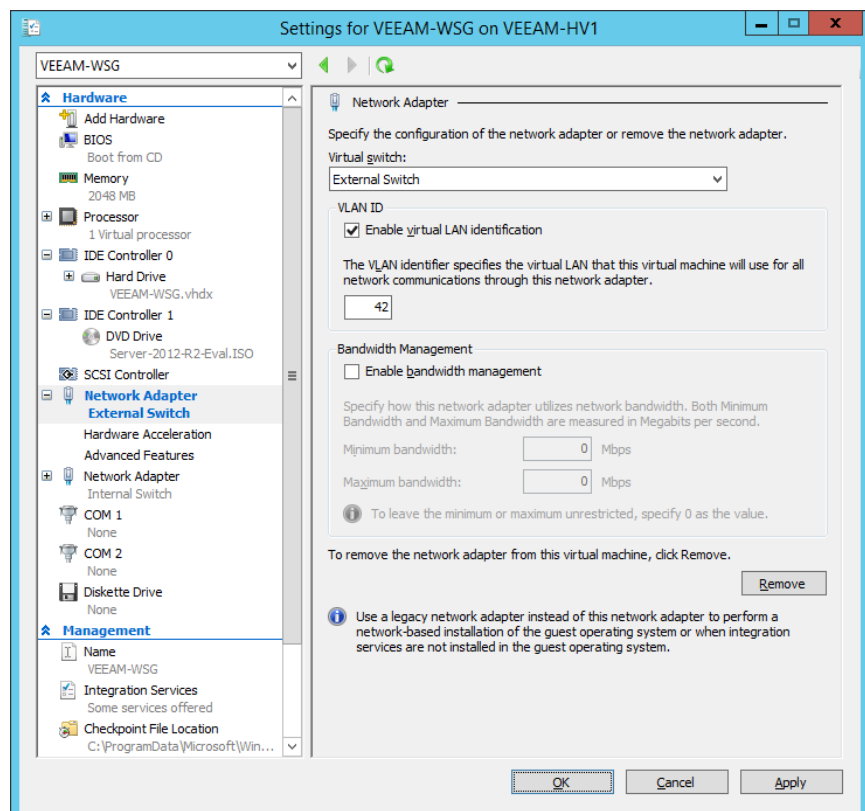


Figure 3g. Configure virtual network adapter VLAN ID

VLAN tags applied at the virtual switch level override VLAN tags applied at the virtual network adapter level. To configure VLAN tags at the virtual switch level, select the **Enable virtual LAN identification for management operating system** option and specify the VLAN identifier. Figure 3h shows the VLAN ID for the External Switch virtual switch set to 2.

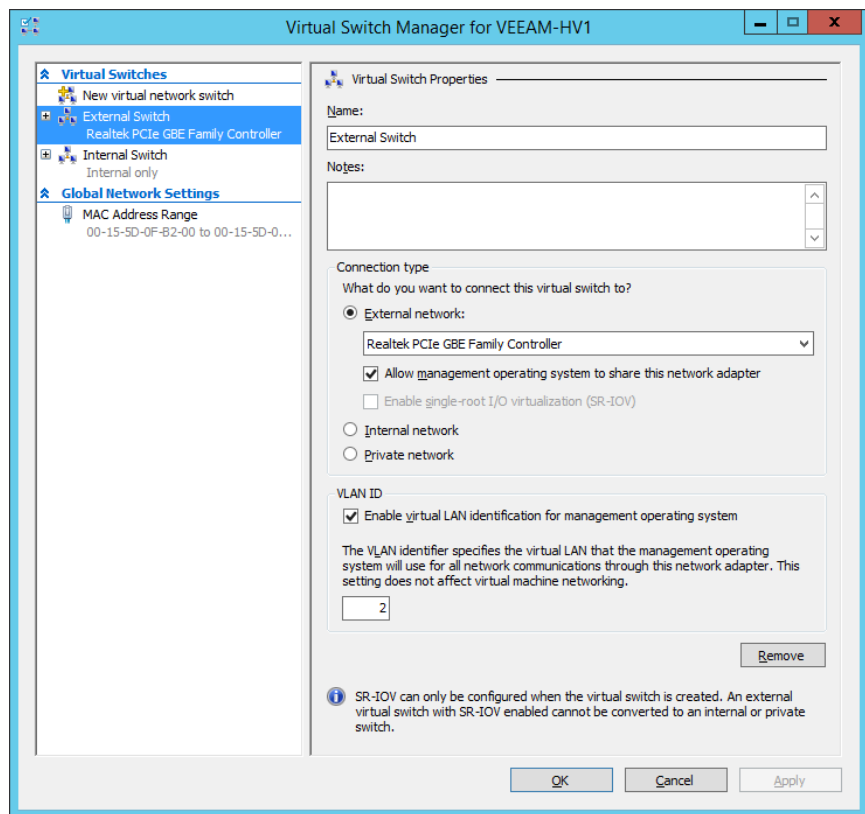


Figure 3h. Configure virtual switch VLAN ID

You'll learn about configuring VLANs and PVLANS using System Center 2012 R2 Virtual Machine Manager (VMM) later in this chapter.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/cc816585%28v=ws.10%29.aspx>

Virtual machine network adapters

Generation 1 VMs support two types of network adapters – synthetic network adapters and legacy network adapters. When you add only one, as shown in the Hyper-V user interface in Figure 3i, it is labelled just Network Adapter.

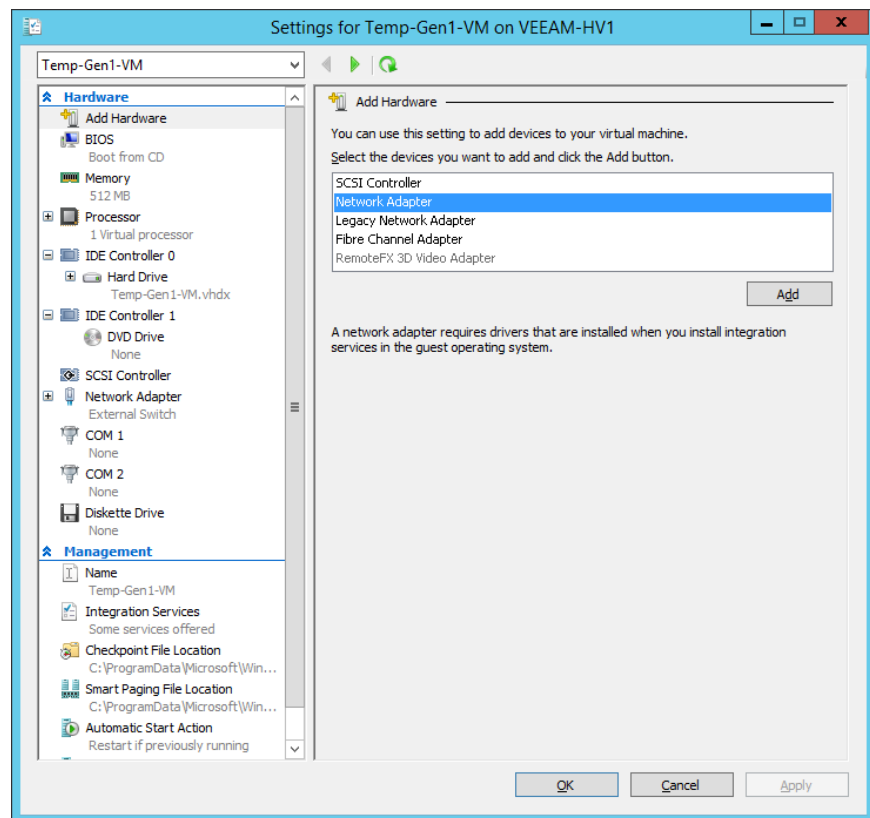


Figure 3i. Network adapter

A synthetic network adapter uses drivers that are provided when you install integration services in the VM operating system. If a VM operating system doesn't have these drivers or if integration services is not available for this operating system, then the network adapter will not function. Synthetic network adapters are unavailable until a VM operating system that supports them is running. This means that you can't perform a PXE boot off a synthetic network adapter if you have configured a generation 1 VM.

Legacy network adapters emulate a physical network adapter, similar to a multiport DEC/Intel 21140 10/100TX 100 MB card. Many operating systems, including those that do not support virtual machine integration services, support this network adapter. This means that if you want to run an operating system in a VM that doesn't have virtual machine integration services support, such as a flavor of Linux that isn't officially supported for Hyper-V, you'll need to use a legacy network adapter as this is likely to be recognized by the guest VM operating system.

Legacy network adapters on generation 1 VMs also function before the VM guest operating system is loaded. This means that if you want to PXE boot a generation 1 VM—for example, if you wanted to use WDS to deploy an operating system to the VM—you'd need to configure the VM with a legacy network adapter.

Generation 2 VMs don't separate synthetic and legacy network adapters and only have a single network adapter type. Generation 2 VMs support PXE booting off this single network adapter type. It is important to remember that only recent Windows client and server operating systems are supported as generation 2 VMs.

Find Out More:

Although this article deals with Hyper-V on Windows Server 2008 R2, you can learn more about the difference between synthetic and legacy network adapters by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/cc770380.aspx>

Virtual machine NIC teaming

NIC teaming allows you to aggregate bandwidth across multiple network adapters while also providing a redundant network connection in the event that one of the adapters in the team fails. NIC teaming allows you to consolidate up to 32 network adapters and to use them as a single network interface. You can configure NIC teams using adapters that are from different manufacturers and that run at different speeds (though it's generally a good idea to use the same adapter make and model in production environments). NIC teaming is a feature in the Windows Server 2012 and Windows Server 2012 R2 operating systems.

You can configure NIC teaming at the virtualization host level if the virtualization host has multiple network adapters. The drawback is that you can't configure NIC teaming at the host level if the network adapters are configured to use SR-IOV. If you want to use SR-IOV and NIC teaming, create the NIC team instead in the VM.

You can configure NIC teaming VMs running on Windows Server 2012 and Windows Server 2012 R2 operating systems by adding adapters to a new team using the Server Manager console. Figure 3j shows the creation of a new NIC team using two virtual network adapters named Veeam-Prod-Team-A.

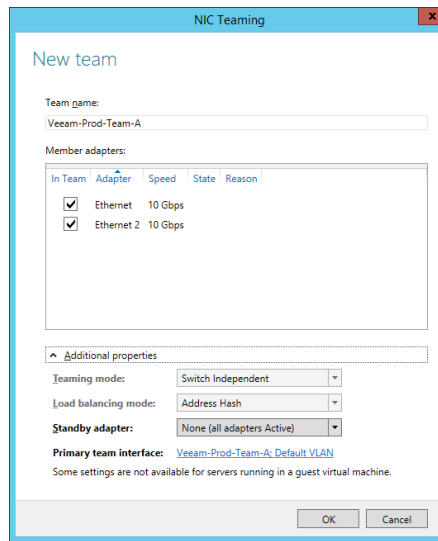


Figure 3j. NIC teaming

When configuring NIC teaming in a VM, ensure that each virtual network adapter that will participate in the team has MAC address spoofing enabled as shown in Figure 3k.

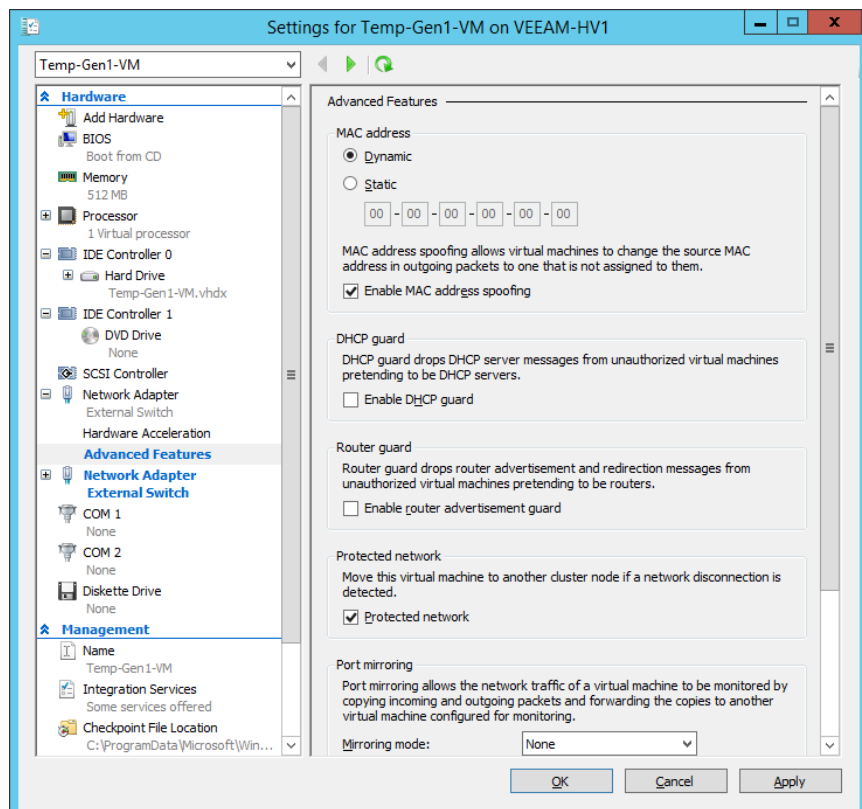


Figure 3k. MAC address spoofing

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/library/XXX.aspx>

VMM logical networks

A logical network in VMM is a collection of network sites, VLAN information, and IP subnet information. You must configure at least one logical network in VMM before you will be able to deploy VMs or services. When you add a Hyper-V virtualization host to VMM, if the virtualization host's physical network adapter is not associated with a logical network, VMM will automatically create one and associate that logical network with the DNS suffix used by the physical network adapter.

Once you create a logical network, you can create network sites. Network sites allow you to associate IP subnets, VLANs and PVLANS with the logical network. Figure 3I shows a network site named Veeam-Logical-Network_0 that has the IP subnet 172.16.42.0/24 associated with it. You can also create IP address pools.

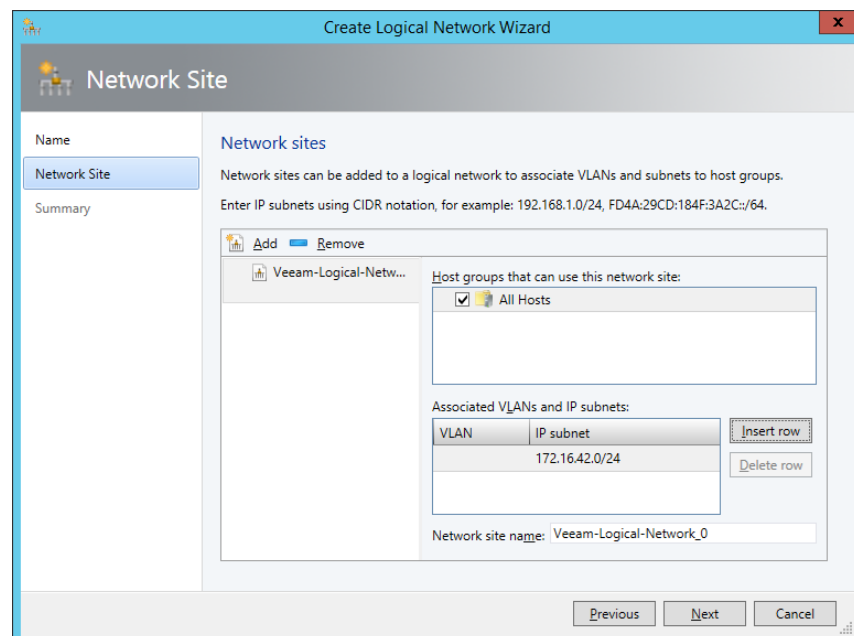


Figure 3I. Create network sites

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/jj721568.aspx>

VMM port profiles and logical switches

VMM port profiles and logical switches function as collections for configuration settings for network adapters across multiple virtualization hosts. Rather than having to configure the properties of each network adapter individually, you apply the configuration information in the logical switches and port profiles. For example, you may want to configure certain network adapters so that they are reserved for use in managing the virtualization host. To accomplish this, you create a logical switch, add the host management port profile to that logical switch, and then assign the logical switch to each adapter that would be used for virtualization host management.

When configuring a VMM logical switch, you configure extensions, uplinks and virtual ports. Extensions allow you to configure what the logical switch does with network traffic. VMM supports the following switch extensions:

- **Monitoring.** Allows you to monitor network traffic. Does not allow you to modify network traffic.
- **Capturing.** Allows you to inspect network traffic. Does not allow you to modify network traffic.
- **Filtering.** You can use filtering extensions to modify, defragment or block packets.
- **Forwarding.** Allows you to modify the destination of traffic based on traffic properties

Uplinks allow you to select which logical networks can connect through a specific physical network adapter. Port classifications allow you to apply functionality-based configurations. For example, they allow you to designate a particular network adapter as being used for VM live migration or network load balancing. VMM includes the following port classifications, which you can add to a virtual switch as shown in Figure 3m.

- **SR-IOV.** Allows a virtual network adapter to use SR-IOV
- **Host management.** For network adapters used to manage the virtualization host
- **Network load balancing.** To be used with network adapters that participate in Microsoft Network Load Balancing.
- **Guest dynamic IP.** Used with network adapters that require guest dynamic IP addresses.

- **Live migration workload.** Used with network adapters that support VM live migration workloads.
- **Medium bandwidth.** Assign to network adapters that need to support medium bandwidth workloads.
- **Host cluster workload.** Assign to network adapters that are used to support host clusters.
- **Low bandwidth.** Assign to network adapters that need to support low-bandwidth workloads.
- **High bandwidth.** Assign to network adapters that are used to support high-bandwidth workloads.
- **iSCSI workload.** Assign to network adapters that are used to connect to SAN resources using the iSCSI protocol.

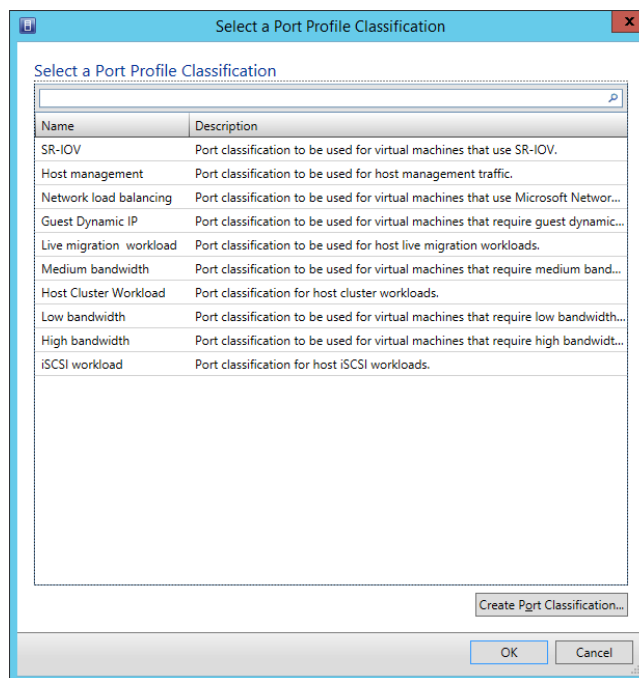


Figure 3m. Port classifications

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/jj721570.aspx>

Network Virtualization

Network virtualization allows you to configure logical networks so that different VM tenants can be assigned the same IP addresses on the same Hyper-V host without collisions occurring. When you configure network virtualization, each VM is assigned two IP addresses – the customer IP address and the provider IP address. The customer IP address is the one that is returned when you examine IP address configuration information from within the VM operating system. The provider IP address is visible in VMM, but is not visible from within the VM operating system. You enable network virtualization when creating a VMM logical network by selecting the checkbox for **Allow new VM networks created on this logical network to use network virtualization** as shown in Figure 3n.

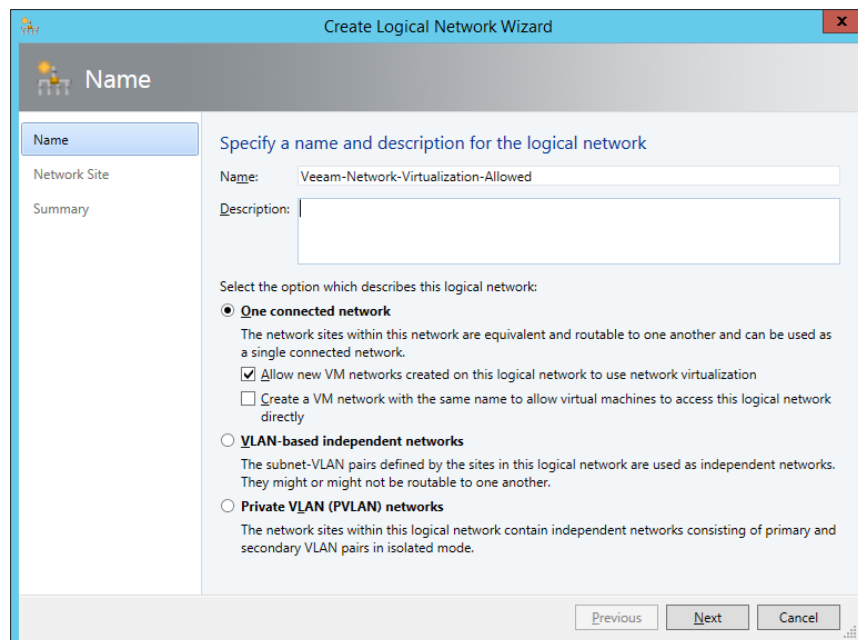


Figure 3n. VM network

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/scvmm/archive/2013/11/27/adopting-network-virtualization-part-ii.aspx>

VMM virtual machine networks

VM networks provide the interface through which VMs connect to a VMM logical network. When you create a logical network, you can have VMM automatically create an associated VMM virtual machine network. Figure 3o shows the VM network named Veeam-VM-Network being associated with the logical network named Veeam-Logical-Network. When configuring a VM network adapter, you connect it to the VM network rather than the configured logical network. If a logical network is configured for network virtualization, you can connect multiple VM networks to the same logical network, but these logical networks will be isolated from one another.

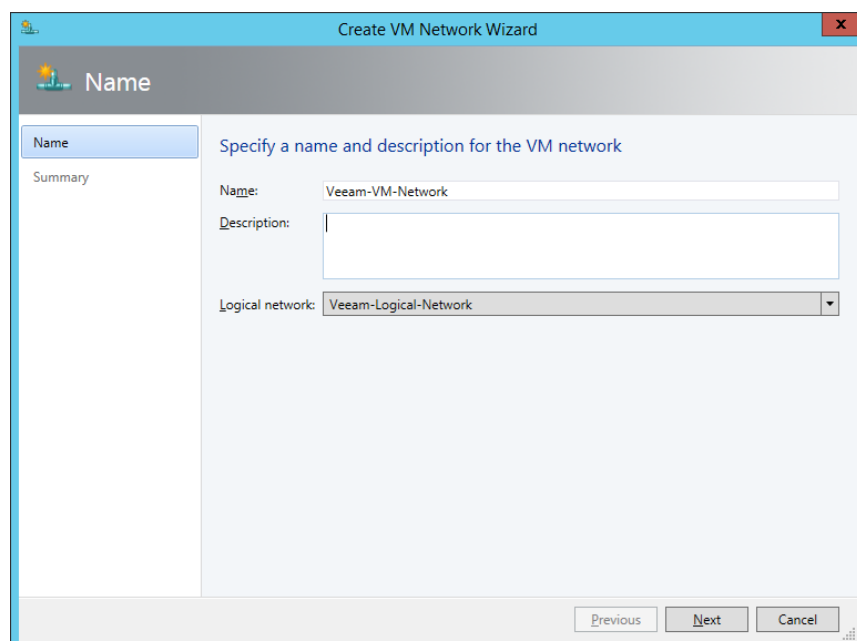


Figure 3o. VM network

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj983727.aspx>

VMM MAC and IP address pools

When managing large numbers of virtualization hosts with VMM, you can configure a MAC address pool for a group of virtualization hosts, rather than configuring the MAC address pool on a per-virtualization host basis. VM network adapters will be allocated MAC addresses from the pool. VMM will ensure that duplicate MAC addresses are not assigned to different VM network adapters. Figure 3p shows the default MAC address pool.

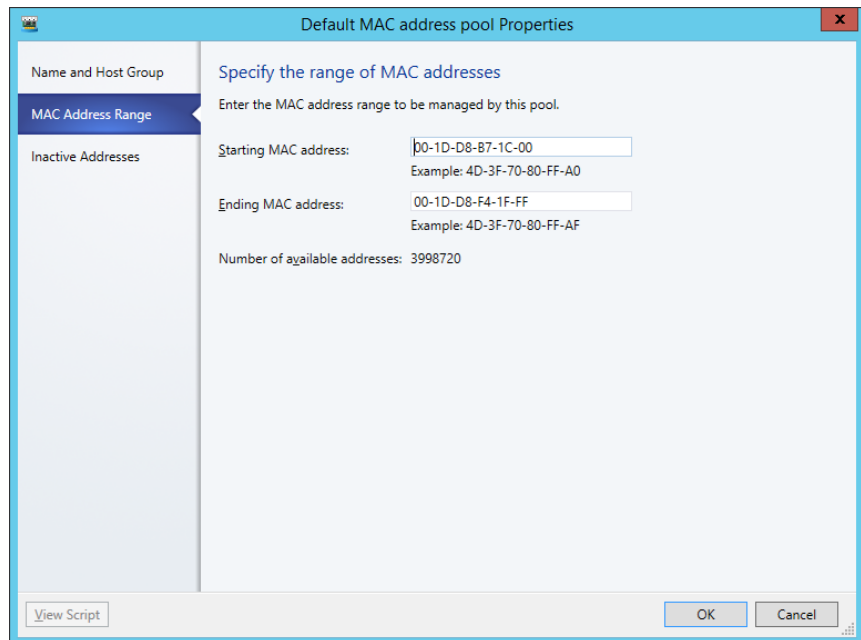


Figure 3p. Configuring MAC address ranges in VMM

Find Out More:

You can learn more about MAC address pools by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610632.aspx>

A static IP address pool is a pool of IP addresses that are associated with an IP subnet that is associated with a network site. For example, if you associate the IP address range 172.16.42.0/24 with a specific network site and then create an IP address pool, the default starting and ending IP addresses for that pool is 172.16.42.1 and 172.16.42.254. Figure 3q shows this range. Static IP address pools also can contain information about default gateway addresses, DNS servers and WINS servers.

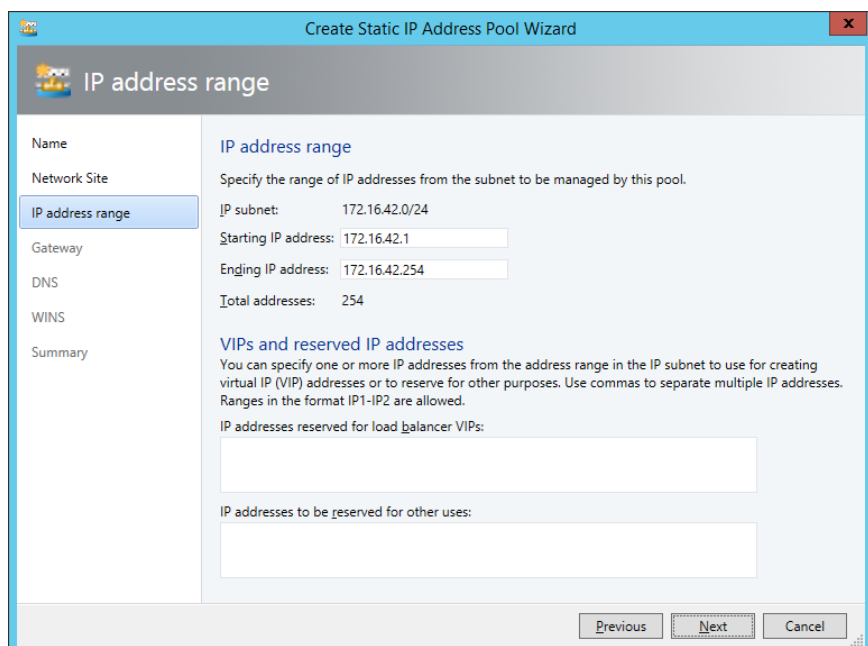


Figure 3q. Configuring IP address pool

VMM can assign VMs running Windows-based operating systems that are associated with this logical network IP addresses from this pool. IP address pools are not mandatory and you can configure VMs to have IP addresses allocated through DHCP. An advantage of IP address pools is that you don't have to worry about creating DHCP reservations to ensure that a specific VM is always assigned the same IP address.

Find Out More:

You can learn more about creating IP address pools by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610590.aspx>

Windows Server Gateway

Windows Server Gateway involves you deploying a specially configured VM that is able to route traffic from private and internal VM networks to one or more external VM networks. The VM you configure as a Windows Server Gateway functions as a virtual router or performs network address translation, depending on your requirements.

When configuring Windows Server Gateway, you need to ensure that the VM that functions as the Windows Server Gateway server will have multiple virtual network adapters. You will need to connect one of these virtual network adapters to an external switch. You will need to connect the other virtual network adapters to internal or private switches. Once the VM has been connected to the appropriate switches and you've installed Windows Server 2012 or Windows Server 2012 R2, you'll need to deploy and configure the Routing role service, which is available as part of the Remote Access role. Figure 3r shows how to configure Routing and Remote Access so that Windows Server 2012 R2 will function as a LAN router.

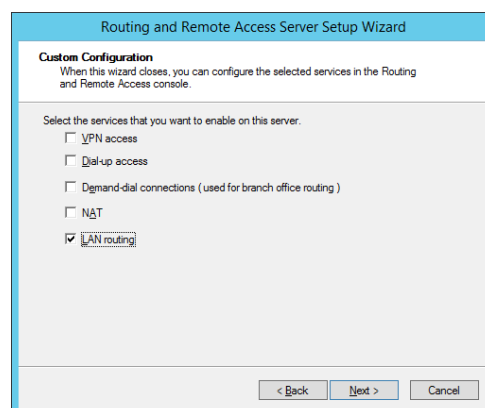


Figure 3r. Configuring MAC address ranges in VMM

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn313101.aspx>

Private Virtual Local Area Networks

VLANs segment traffic at layer 2 of the OSI model by tagging packets. As you learned earlier, a VLAN ID is a 12-bit number, allowing you to allocate VLAN IDs between 1 and 4095. In very large data centers that host a large number of tenants, you may need more VLANs than are available with a 12-bit address. Private Virtual Local Area Networks (PVLANS) allow you to separate a VLAN into multiple isolated sub-networks. When you implement PVLANS, there are two VLAN identifiers, the primary VLAN ID and the secondary VLAN ID, as shown in Figure 3s.

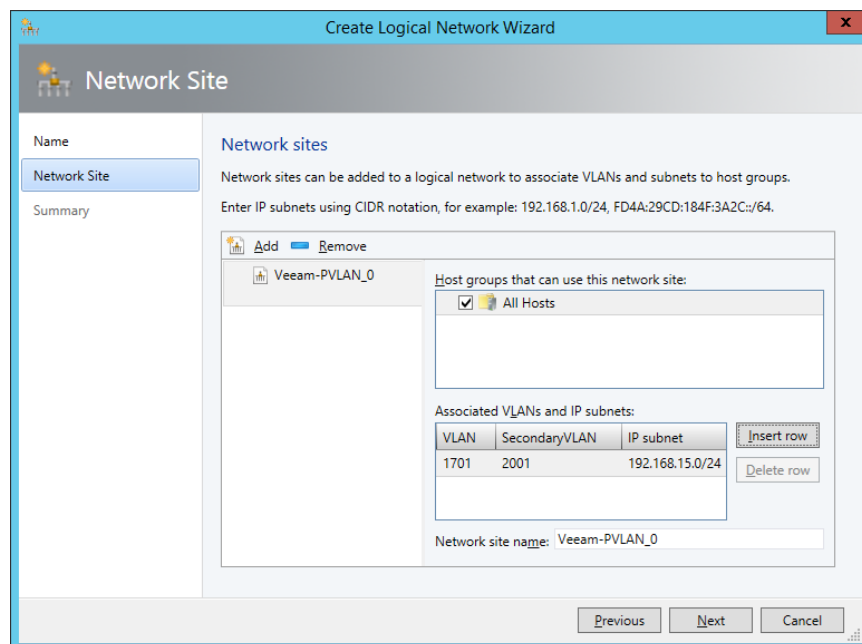


Figure 3s. Configuring PVLAN in VMM

You implement VLANs and PVLANS in VMM by creating logical networks. When you create the logical network, you specify the VLAN and/or PVLAN ID as well as the IPv4 or IPv6 network.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: http://technet.microsoft.com/en-us/library/jj679878.aspx#bkmk_pvlan

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. What type of Hyper-V virtual switch would you configure if you wanted to allow network communication between the Hyper-V host and a VM, but not allow the VM access to any external network?
2. How can you ensure that a particular VM on a Hyper-V host never uses more than 50 Mbit/sec bandwidth?
3. You want to ensure that a particular VM is assigned a specific address from a DHCP reservation. What step would you need to take to accomplish this goal?
4. You have three VMs connected to the same external virtual switch. You want to ensure that the traffic from these VMs is in separate broadcast domains. How can you accomplish this goal?
5. You want to PXE boot a generation 1 VM. What type of network adapter should you configure to accomplish this goal?
6. Which Windows Server operating systems include NIC teaming as a feature?
7. What configuration change must you make for VM network adapters if you are going to create a NIC team within the VM guest operating system?
8. What do you need to configure if you want VMM to assign IP addresses to VMs without using a DHCP server?
9. Which role service do you need to install on a VM running the Windows Server 2012 R2 operating system if you are going to configure it to function in the Windows Server Gateway role?
10. You need to support more than 5,000 tenants in your organization's virtualization deployment. You manage this deployment using System Center 2012 R2 VMM. You want to ensure that traffic from each tenant is isolated from every other tenant. Which option would you choose when creating a logical network in VMM to accomplish this goal?

Chapter 4: Implementing Virtual Machines

Deploying a single virtual machine (VM) to a single Hyper-V server is a relatively straightforward task. You answer a few questions in a wizard and soon after, the VM is deployed. Deploying large numbers of VMs across large numbers of VM hosts is more complicated. In this scenario you need to take into account issues around high availability, selecting the best host for the VM, and simplifying the deployment of large numbers of similarly configured VMs through the utilization of templates.

In this chapter you'll learn about:

- Highly available VMs
- Guest resource optimization
- Placement rules
- VMM templates

Highly available virtual machines

Highly available VMs remain available even when the virtualization host that is hosting the VM suffers a failure. There are two general methods of making VMs highly available. The first, and most traditional method, is to deploy the VM on a Hyper-V failover cluster. Deploying a VM on a Hyper-V failover cluster has the benefit of ensuring that there is minimal, if any, disruption for clients connected to the VM if a Hyper-V host fails. Chapter 6, "Hyper-V Failover Clustering and Failover Clustering Roles," covers highly available VMs in more detail.

The second method of making VMs highly available is to configure Hyper-V replication. Hyper-V replication provides a replica of a VM running on one Hyper-V host that can be stored and updated on another Hyper-V host. For example, you could configure a VM hosted on a Hyper-V failover cluster in Melbourne to be replicated through Hyper-V replica to a Hyper-V failover cluster in Sydney. Hyper-V replication allows for replication across site boundaries and does not require access to shared storage in the way that failover clustering does.

Hyper-V replication is asynchronous. While the replica copy is consistent, it is a lagged copy with changes sent only as frequently as once every 30 seconds. Hyper-V replication supports multiple recovery points, with a recovery snapshot taken every hour (this incurs a resource penalty, so is off by default). This means that when activating the replica, you can choose to activate the most up-to-date copy or a lagged copy. You would choose to activate a lagged copy in the event that some form of corruption or change made the up-to-date copy problematic. Figure 4a shows configuring additional recovery points when enabling Hyper-V replica.

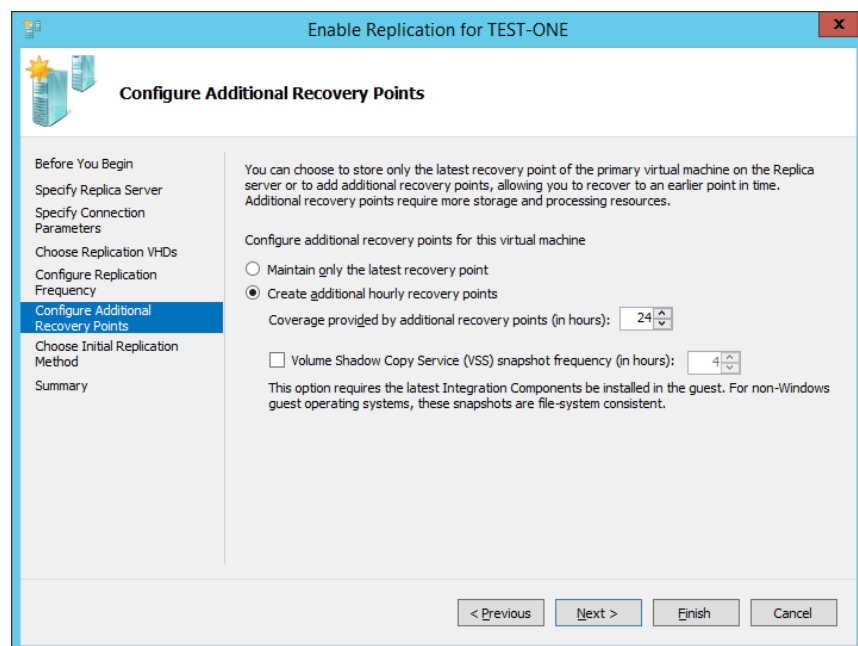


Figure 4a. Configuring additional recovery points

When you perform a planned failover from the primary host to the replica, you need to switch off the primary. This ensures that the replica is in an up-to-date consistent state. This is a drawback compared to failover or live migration where the VM will remain available during the process. A series of checks are completed before performing planned failover to ensure that the VM is off, that reverse replication is allowed back to the original primary Hyper-V host and that the state of the VM on the current replica is consistent with the state of the VM on the current primary. Performing a planned failover will start the replicated VM on the original replica, which will now become the new primary server.

Hyper-V replica also supports unplanned failover. You perform an unplanned failover in the event that the original Hyper-V host has failed or the site that hosts the primary replica has become unavailable. When performing unplanned failover, you can either choose the most recent recovery point or choose a previous recovery point. Performing unplanned failover will start the VM on the original replica, which will now become the new primary server.

Windows Server 2012 R2 supports Hyper-V extended replication. Hyper-V extended replication allows you to create a second replica of the existing replica server. For example, you could configure Hyper-V replication between a Hyper-V virtualization host in Melbourne and Sydney, with Sydney hosting the replica. You could then configure an extended replica in Brisbane using the Sydney replica

To configure Hyper-V replica, you need to enable each virtualization host that will participate in the replica as a replica server. Doing this involves specifying an authentication method—either Kerberos or certificate-based authentication. You also need to specify whether the replica will accept VMs from any replicated server or only specific servers. Figure 4b shows a Hyper-V server on which Hyper-V replica has been enabled.

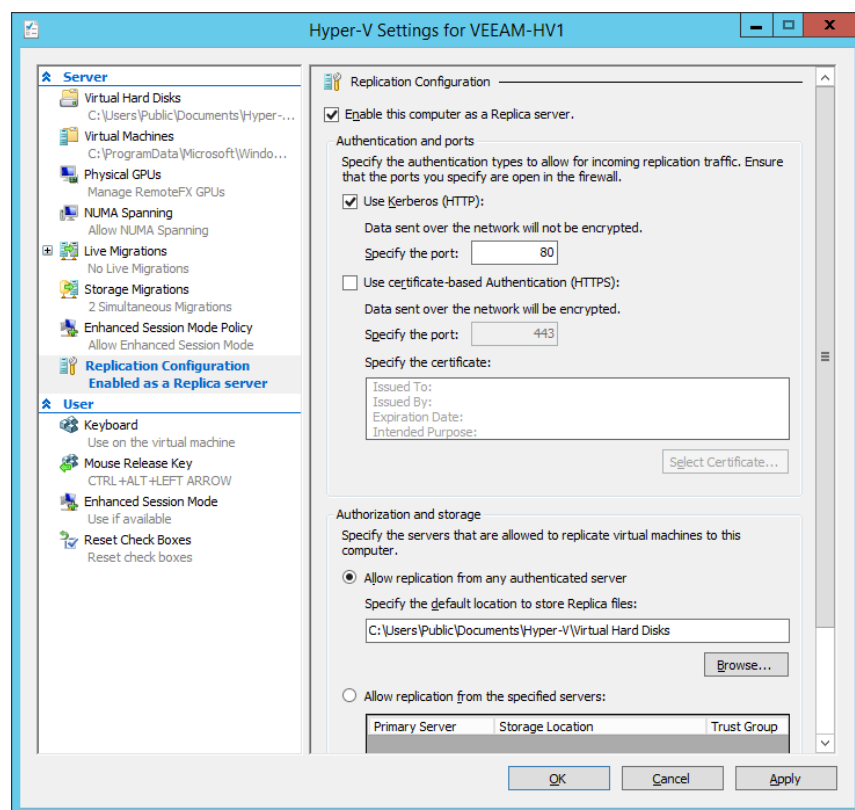


Figure 4b. Configuring Hyper-V replication

Find Out More:

You can learn more about Hyper-V replica, consult the following TechNet article: <http://technet.microsoft.com/en-us/library/jj134172.aspx>

Guest resource optimization

When you are deploying a VM and you have the option of deploying to different virtualization hosts, you generally want to ensure that you deploy the VM to a virtualization host that has the greatest amount of available resources. Once deployed, you also want to ensure that your virtualization hosts are being used equitably and that VMs are not being hosted disproportionately on some virtualization hosts while the resources on other virtualization hosts remain idle.

Intelligent placement

When you are preparing to deploy or migrate a VM to more than one virtualization host, VMM will use intelligent placement to assess the suitability of each virtualization host. This assessment is provided as a rating. Ratings are based on CPU, memory, network and disk resource utilization. Figure 4c shows the ratings for Hyper-V virtualization hosts veeam-hv1.veeam.internal and veeam-hv2.veeam.internal.

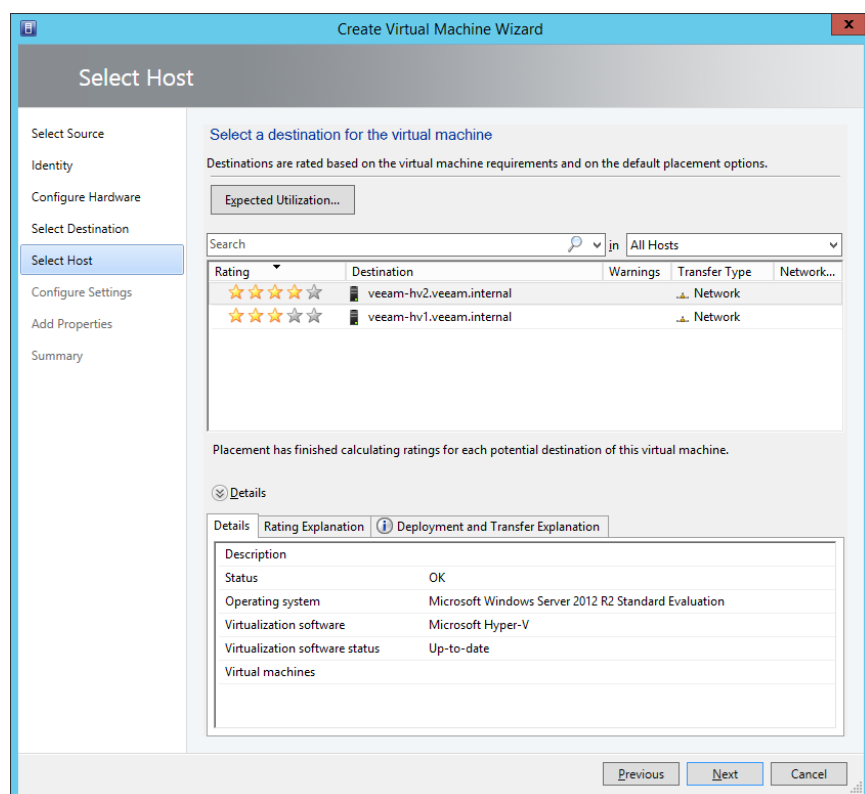


Figure 4c. Reviewing intelligent placement ratings

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj860428.aspx>

Dynamic optimization

Dynamic optimization allows the automatic migration of VMs within a VMM managed virtualization host cluster to balance VM load across cluster nodes. Dynamic optimization can only be used to migrate VMs that are hosted on shared storage. VMs that are not configured as highly available cannot be migrated using dynamic optimization.

You can configure dynamic optimization for clusters with two or more nodes. While you can configure dynamic optimization at the host group level, any host that is a member of the host group that isn't a member of a cluster, or any cluster that does not support live migration will not participate in the dynamic optimization process.

When configuring dynamic optimization, you specify how aggressive the process should be. This determines whether migrations will be initiated only for major performance gains or, if set to a higher aggressiveness level, for minor gains as well. You can configure how often this assessment is performed, with the default being every 10 minutes. You can also specify the threshold at which hosts will be eligible for optimization. Figure 4d shows the configuration of dynamic optimization.

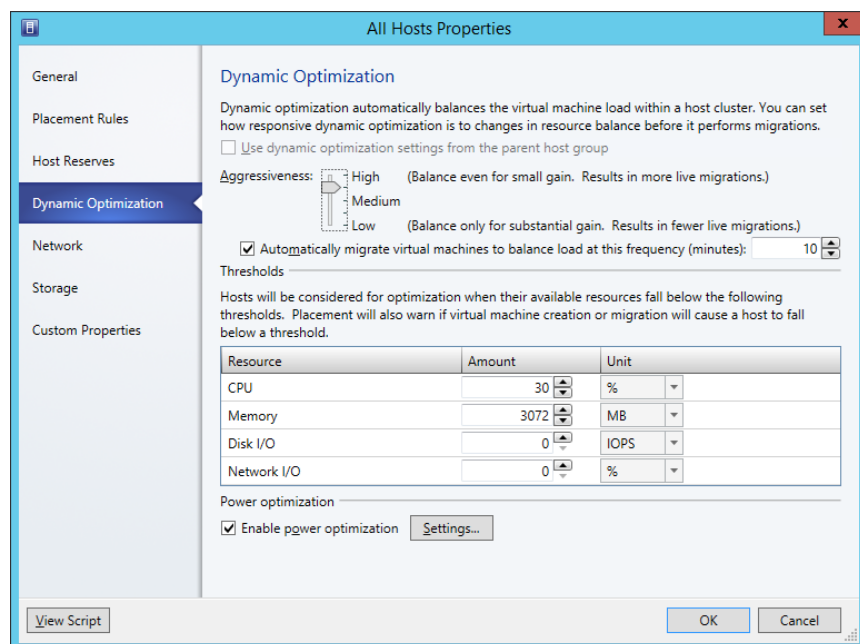


Figure 4d. Dynamic optimization settings

Power optimization

Power optimization allows VMM to power off virtualization hosts that are not required; for example, shutting down one Hyper-V host in an eight-node cluster when seven nodes are more than adequate to host the current VM workload. Power optimization also allows virtualization hosts to be powered back on when workload requirements increase and the extra host capacity becomes necessary. You can only use power optimization when you configure a VMM host group to allow VM migration to be driven by dynamic optimization.

When configuring power optimization, you specify the threshold value, in terms of the performance of other hosts in the cluster, that must be obtained in the event that the server selected for shut down has all of its currently running VMs live migrated away. Figure 4e shows the default configuration settings for power optimization.

Customize Power Optimization Schedule

Power optimization settings

Power optimization will try to evacuate hosts of a balanced cluster and turn them off to save power.

Thresholds

Hosts will be considered for power optimization if they can be evacuated without causing any remaining nodes of the cluster to fall below the following thresholds.

Resource	Amount	Unit
CPU	40	%
Memory	4096	MB
Disk I/O	0	IOPS
Network I/O	0	%

Schedule

Select the days and times when you want power optimization to run. Times are applied locally to the time zone of each virtualization host.

	Midnight (AM)												Noon (PM)												
	12	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11	12
Sun	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Mon	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Tue	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Wed	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Thu	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Fri	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Sat	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

No power optimization. Any hosts that were shut down by dynamic optimization are restarted.
 Power optimization is running. Hosts are shut down and restarted as needed.

OK Cancel

Figure 4e. Power optimization settings

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/gg675109.aspx>

Performance and Resource Optimization

Performance and Resource Optimization (PRO) allows you to use System Center 2012 R2 Operations Manager to monitor virtualization hosts that are hosted by System Center 2012 R2 Virtual Machine Manager (VMM). When configured with a PRO-enabled management pack, an alert raised in Operations Manager, such as excessive CPU, storage, or memory utilization on a Hyper-V host can create a PRO tip in VMM. A VMM PRO tip can perform an action to resolve the operations manager alert, such as live migrating VMs to different nodes in a Hyper-V cluster, or to live migrate to a separate Hyper-V cluster.

You enable PRO when configuring integration with the Operations Manager server, as shown in Figure 4f. You may need to disable dynamic optimization and power optimization or reconfigure settings if you are using PRO-enabled management packs.

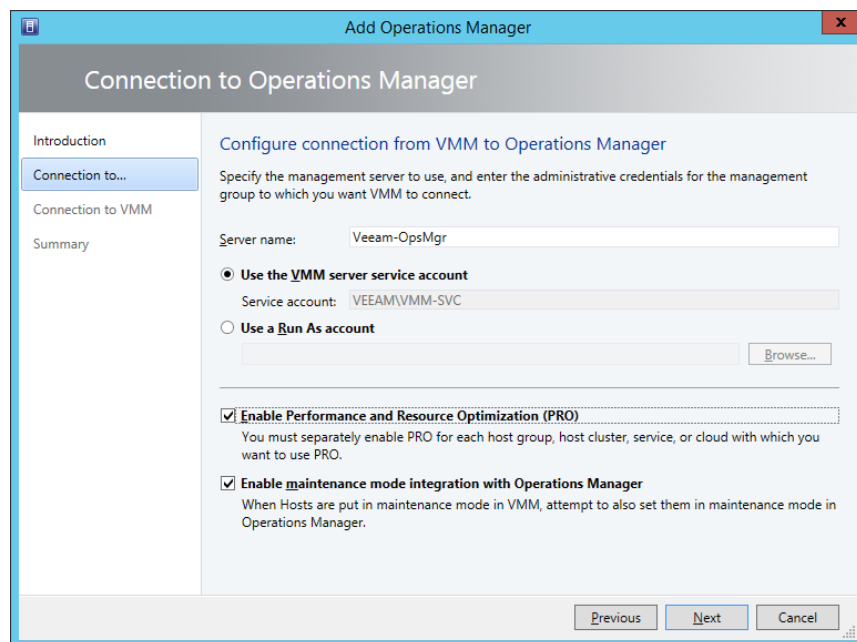


Figure 4f. Enabling Performance and Resource Optimization

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
<http://technet.microsoft.com/en-us/library/gg675109.aspx>

Placement rules

Placement rules allow you to configure additional rules based on custom properties that influence where VMM will place a VM. The first step to take in configuring placement rules is to create custom properties. You can do this by selecting **Custom Properties** on the Host Group properties dialog. Figure 4g shows the creation of two custom properties, Department and Branch Location.

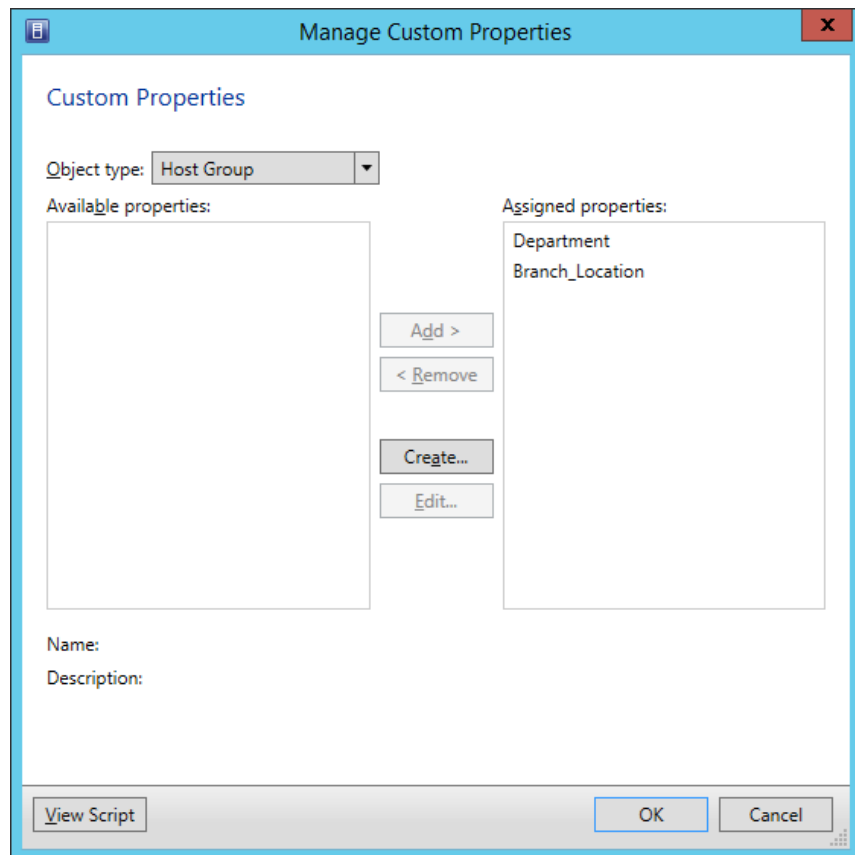


Figure 4g. Configuring custom properties

Once you have custom properties, you can assign them to VMs, hosts, host groups and templates. You can then create rules using operators that will influence VM placement (see Figure 4h). The operators have the following influence on placement:

- **Must match.** 0 stars and the VM will be blocked if the custom property of the VM doesn't match the custom property of the host.
- **Should match.** A placement warning will be issued if the custom property of the VM doesn't match the custom property of the host.
- **Must not match.** 0 stars and the VM will be blocked if the custom property of the VM does match the custom property of the host.
- **Should not match.** A placement warning will be issued if the custom property of the VM does match the custom property of the host.

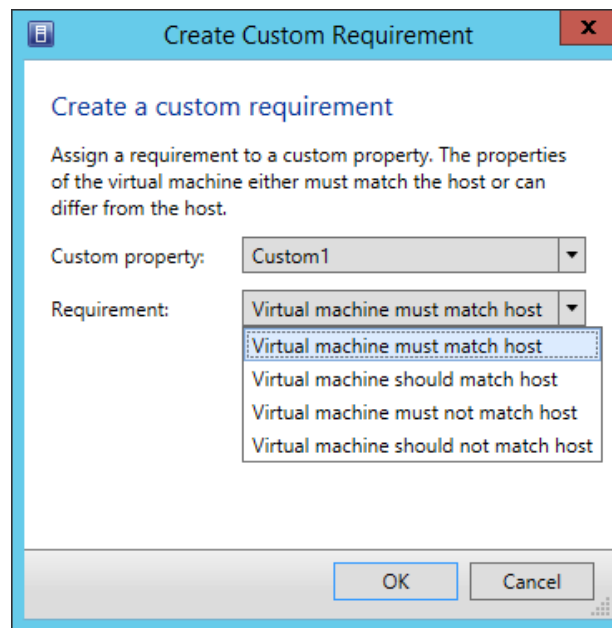


Figure 4h. Configuring placement rules

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/scvmm/archive/2013/03/11/custom-placement-rules-and-availability-sets-in-scvm-2012-sp1.aspx>

VMM templates

VMM virtual machine templates allow you to create VMs that have a consistent set of configuration settings. A VMM virtual machine template is an XML object stored in VMM libraries. Templates have one or more of the following components associated with them:

- **Guest OS profile.** Provides configuration information about the operating system that will be deployed to the VM.
- **Hardware profile.** Hardware profiles specify hardware configuration settings including CPU, memory and network adapter configuration.
- **Virtual hard disks.** This can be a blank or a specially prepared virtual hard disk; for example, a virtual hard disk that has a sysprepped version of Windows Server 2012 R2 on it.

Before you create a VMM template, you should create a guest OS profile and a hardware profile. You don't have to do this, but it saves you from having to create these profiles when creating the template. You can also configure application profiles, capability profiles, physical computer profiles and SQL server profiles. These allow you to configure SQL Server deployment, or deployment of another application on a host as part of deploying a template. SQL Server profiles require a specially prepared virtual hard disk with SQL Server installed. Physical computer profiles are used when using VMM to deploy an operating system to a computer that will function as a virtualization host.

Guest OS profile

A guest OS profile includes operating system settings that you want to apply to the VM during deployment. You can configure the following settings when configuring a Windows guest OS profile:

- **Operating system.** If you select a Windows operating system, you can select any Windows operating system and edition, from Windows 2000 Advanced Server through to Windows Server 2012 R2 Datacenter edition.
- **Identity information.** Allows you to set the computer name assigned to the VM by using a combination of characters and wildcards.
- **Admin password.** Use this to configure a local administrator password. Other options include configuring no local administrator password or specifying a Run As account for the local administrator account.
- **Product key.** Allows you to enter a product key, usually a Multiple Activation Key (MAK) to be assigned to the VM. Note that if the Hyper-V virtualization host is running Windows Server 2012 R2 Datacenter edition, you'll be able to use automatic activation for supported VMs.
- **Time zone.** Allows you to specify the time zone that is assigned to the VM.
- **Roles.** Use this to specify which roles should be installed on the computer. This varies depending on which operating system is selected. Figure 4i shows some of the roles available for Windows Server 2012 R2.

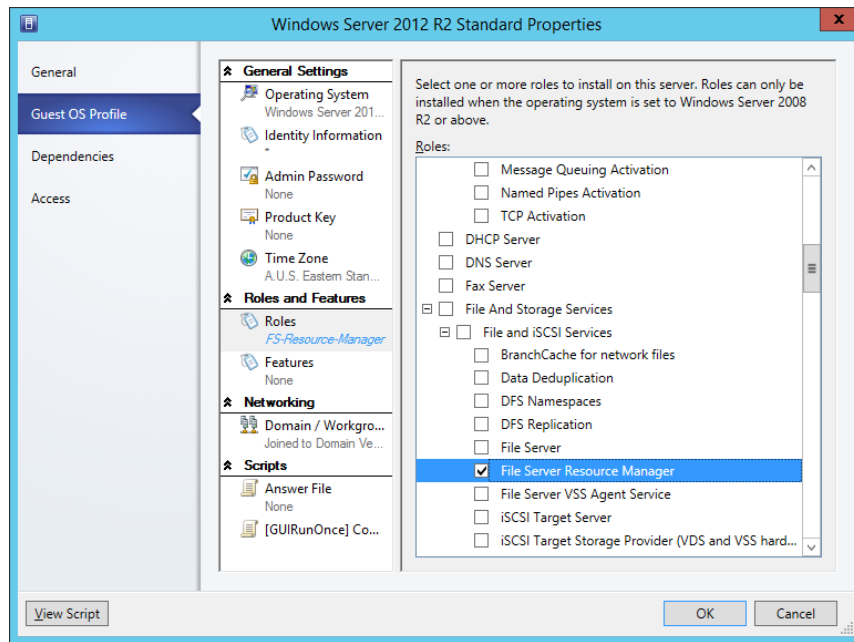


Figure 4i. Roles in guest OS profile

- **Features.** Use this to specify which features should be installed on this computer. This will vary depending on which operating system you select.
- **Domain/Workgroup.** Allows you to configure whether the computer will be assigned to a workgroup or joined to a domain. Figure 4j shows the domain set to Veeam.internal and also shows where the domain credentials are used to join the computer to the domain.

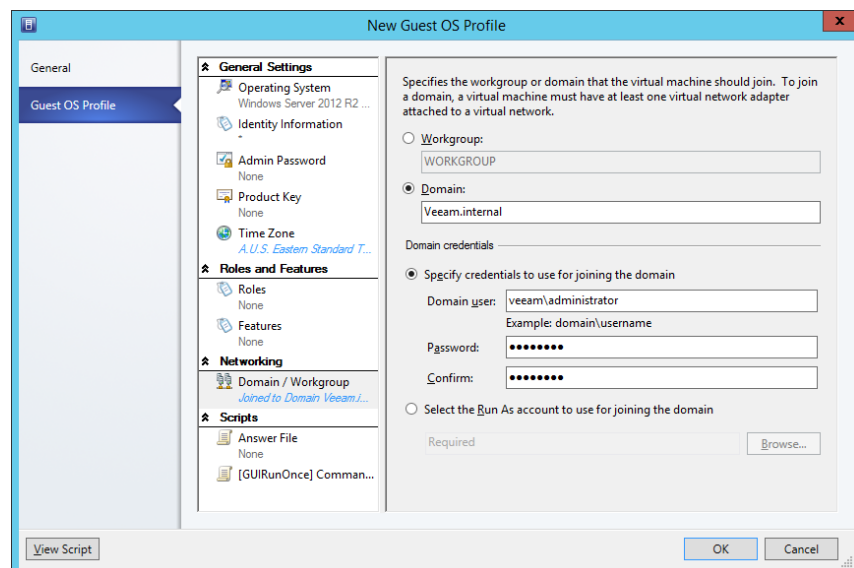


Figure 4j. Guest OS profile domain join

- **Answer file.** Allows you to specify an answer file. For Windows Server 2003 this will be "Sysprep.inf." For Windows Server 2008 and later this will be "unattend.xml." You can generate this file using the Windows System Image Manager (SIM) tool available in the Windows ADK.
- **GUIRunOnce command.** Use this to specify a command that will run the first time a user signs on to the computer.

When configuring a Linux Guest OS Profile, you can configure the following settings as shown in Figure 4k:

- **Linux OS.** VMM guest OS profiles can be configured for supported versions of CentOS Linux, Debian GNU/Linux, Oracle Linux, Red Hat Enterprise Linux, Suse Linux Enterprise Server and Ubuntu Linux.
- **Identity information.** Allows you to specify the VM computer name.
- **Root credentials.** Allows you to specify root account credentials as well as SSH key information.
- **Time zone.** Allows you to specify the time zone
- **Scripts.** Allows you to specify scripts that should be run to configure the operating system during deployment.

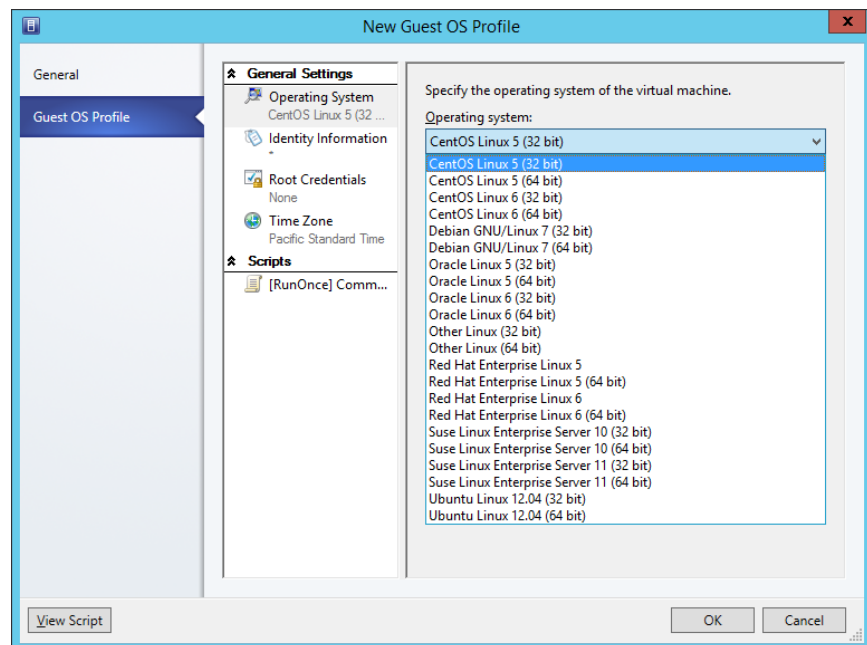


Figure 4k. Linux guest OS profile

Find Out More:

You can learn more about Guest OS profiles by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh427296.aspx>

Hardware profile

VMM hardware profiles allow you to specify the virtual hardware configuration of VMs. You can configure the following options when creating a hardware profile:

- **Generation.** Allows you to specify the VM generation. Hyper-V virtualization hosts running Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 and Windows Server 2012 R2 support generation 1 VMs. Only Hyper-V virtualization hosts running Windows Server 2012 R2 support generation 2 VMs.
- **Cloud compatibility.** Use this to specify the capability profile for the virtualization host that the VM will be deployed on. You can choose Hyper-V, XenServer or ESX Server. Figure 4l shows Hyper-V selected.

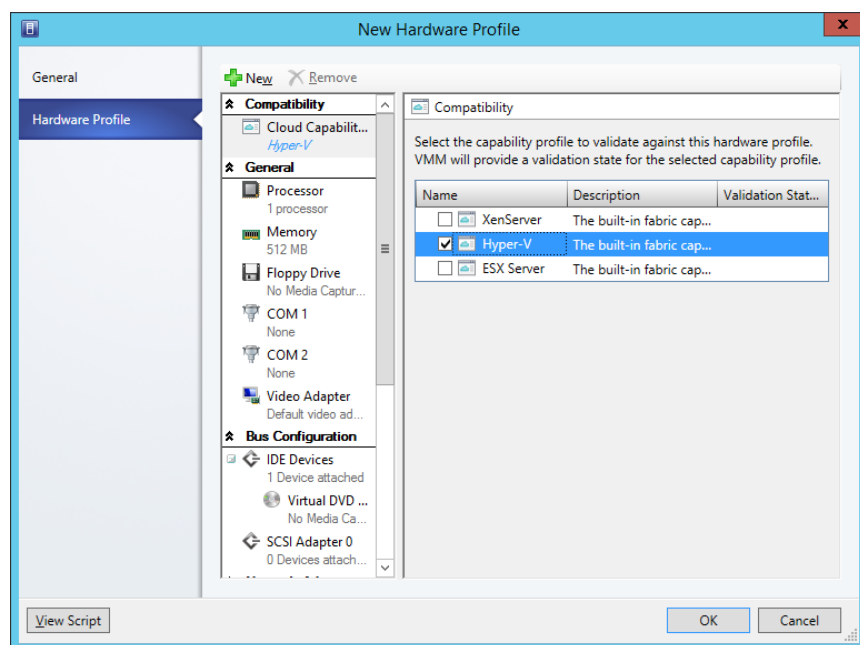


Figure 4l. Cloud compatibility

- **Processor.** Use this setting to specify the number of processor cores to be allocated to the VM. You can also use this setting to enable compatibility mode, allowing migration of the VM to a virtualization host that uses a different processor version.
- **Memory.** Specify the amount of memory to be assigned to the VM, as shown in Figure 4m. You can configure static or dynamic memory using this setting.

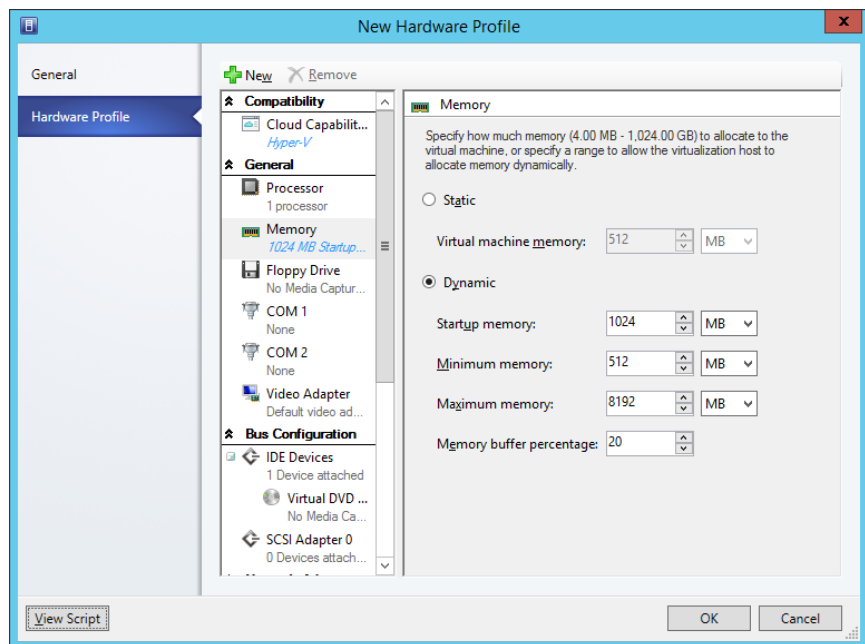


Figure 4m. Memory configuration

- **Floppy drive.** Allows you to mount a floppy image (only supported on generation 1 VMs).
- **Com 1.** Allows you to configure a named pipe or physical COM port to a virtual COM 1 on a generation 1 VM.
- **Com 2.** Allows you to configure a named pipe or physical COM port to a virtual COM 1 on a generation 1 VM.
- **Video adapter.** Use this to specify whether a standard video adapter or a Microsoft RemoteFX 3D video adapter will be configured for the VM. Remote FX 3D video adapters are not supported on generation 2 VMs.
- **IDE devices.** Use this to configure which virtual devices are connected to the virtual IDE adapter (only supported on generation 1 VMs).
- **SCSI adapter 0.** Use this to configure the virtual devices attached to SCSI adapter 0.
- **Network adapter.** Allows you to specify network adapter settings including which VM network the adapter is connected to, whether the adapter uses a dynamic or static IP, whether a dynamic or static MAC address is used as well as port profile settings, including classification, virtual switch optimizations and MAC address spoofing. Figure 4n shows Network Adapter settings for Network Adapter 1.

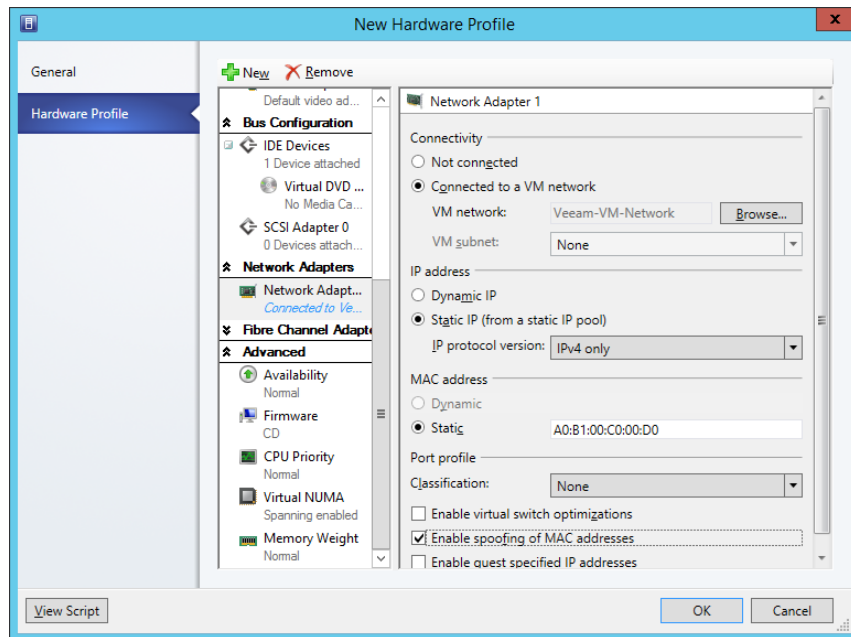


Figure 4n. Network adapters

- **Availability.** Use this setting to specify whether the VM will be highly available. When you do this, VMM will place the VM on a virtualization host that is part of a host cluster. As Figure 4o shows, you can also specify the VM priority, which is used to determine the startup order on the node should a restart occur without the VM being live migrated off the node.

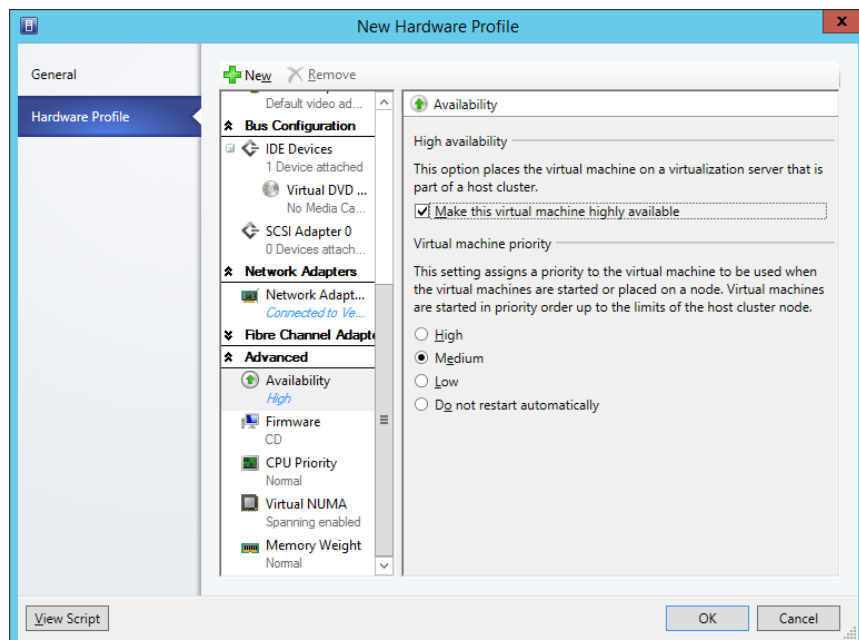


Figure 4o. Availability configuration

- **Firmware.** Allows you to configure the startup order for the VM. By default, the VM will attempt to boot a virtual CD device before attempting an IDE hard drive (generation 1 VM), SCSI hard drive (generation 2 VM, PXE boot, or floppy image (generation 1 VM)).
- **CPU priority.** Allows you to specify a priority CPU use when allocating resources on the virtualization host. Figure 4p shows this setting.

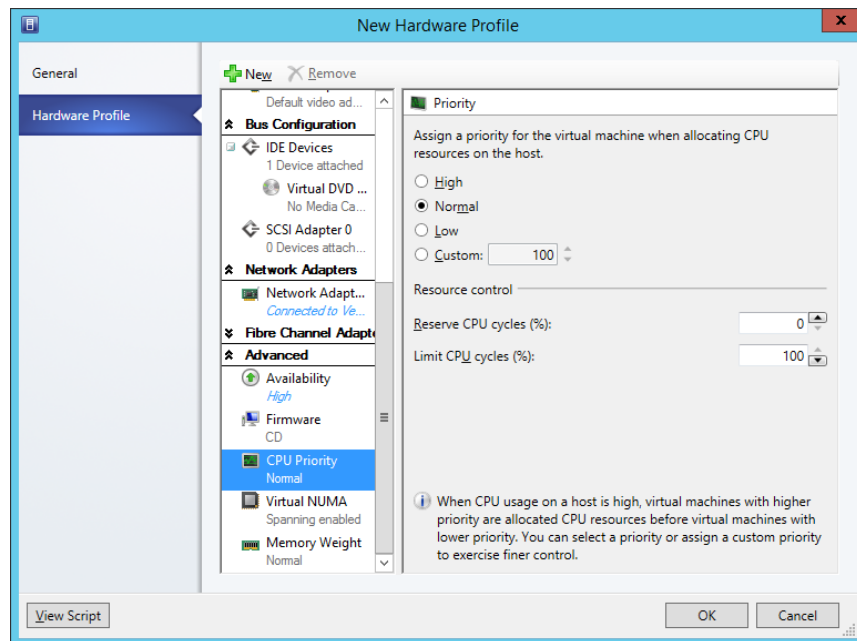


Figure 4p. CPU priority

- **Virtual NUMA.** Allows you to configure Non Uniform Memory Access settings.
- **Memory weight.** Use this to configure the priority given by the virtualization host to the VM's requests for memory.

You can also add the following additional virtual devices to a hardware profile:

- SCSI adapter
- DVD
- Network adapter
- Legacy network adapter
- Fibre channel adapter

Find Out More:

You can learn more about Hardware Profiles by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh427289.aspx>

VMM template configuration

To create a VMM template, you'll need to provide the following information:

- Specify an existing VM on which to base the template, an existing VM template or specify a virtual hard disk that is already stored in the VMM library. Figure 4q shows selection of a virtual hard disk with a sysprepped version of Windows Server 2012 R2 installed on it, which is stored in the library.

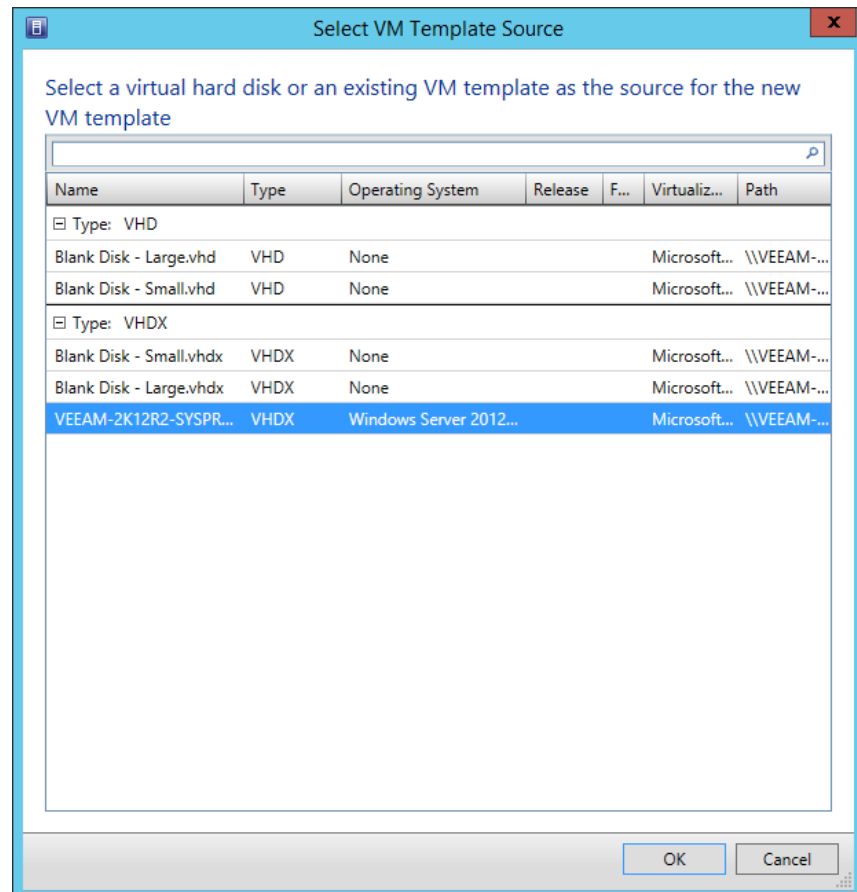


Figure 4q. VMM template source

- Provide a template name and specify whether the VM will be generation 1 or generation 2.
- Select an existing hardware profile or create a new hardware profile.
- Create a new OS profile or select an existing OS profile.
- Create an application profile, or select an existing application profile or choose not to use an application profile.
- Create a SQL Server configuration profile, select an existing SQL Server configuration profile or choose not to use a SQL Server configuration profile.

Find Out More:

You can learn more about VMM Templates by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh427282.aspx>

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. You have configured Hyper-V replica to replicate a VM between Melbourne and Sydney. What state must the VM be in before you can perform a failover from the primary to the replica server?
2. Which VMM feature should you configure if you want an underutilized Hyper-V server to be powered down automatically by VMM?
3. Which other member of the System Center suite must you configure and deploy before you can use Performance and Resource Optimization?
4. Which VMM component allows you to specify the built-in administrator account password, computer name, as well as the roles and features that should be installed?
5. Which VMM component do you configure to specify the number of virtual processors and virtual network adapters a VM will be configured with when deployed?
6. List the three different sources you can use to create a VMM template.

Chapter 5: Managing Virtualization Hosts and Infrastructure

Depending on the size of your organization, you may have a large number of virtualization hosts. When you have a large number of virtualization hosts, you'll need to come up with a way to efficiently manage them. This will involve delegating administrative privileges to members of the IT team, ensuring that the System Center Virtual Machine Manager (VMM) libraries are properly maintained, integrating third-party virtualization platforms, ensuring that virtual machine (VM) images are serviced and that other System Center 2012 R2 components, such as Operations Manager, Data Protection Manager and Service Manager are correctly integrated.

In this chapter you'll learn about:

- Delegating virtualization management tasks
- Managing VMM libraries
- Integrating third-party virtualization platforms
- Deploying Hyper-V hosts to bare metal
- Integrating Operations Manager and Service Manager
- Servicing VM images
- Integrating Data Protection Manager

Delegating virtualization management tasks

Role-Based Access Control (RBAC) is a principle related to the delegation of privileges. It involves both delegating privileges and specifying a scope for that delegation. For example, rather than just delegate a privilege, such as the ability to deploy a VM, across the entire virtualization environment, you could instead delegate the ability to deploy VMs from specific templates only to specific virtualization hosts.

VMM implements RBAC through user roles. You create user roles by selecting role profiles, members, a scope, library servers and Run As accounts, as shown in Figure 5a. Depending on the role profile selected, you can also configure a user role to contain networks, cloud quotas, resources and permissions.

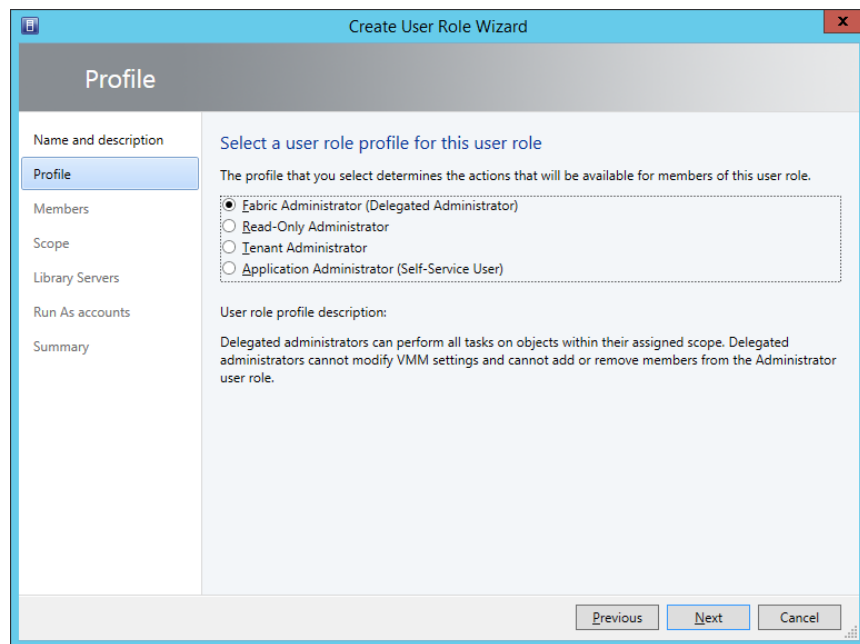


Figure 5a. Select role profile

Role Profiles

The role profile you choose determines what actions an account that is assigned a user role can perform. VMM includes the following role profiles:

- **Administrator.** Members of this role are able to perform any administrative action on any object that VMM can manage. Administrators are able to add XenServer hosts and clusters to VMM. Administrators are also able to add Windows Server Update Services (WSUS) to VMM to facilitate the management of software updates.
- **Fabric Administrator (Delegated Administrator).** Members of this role are able to perform all administrative tasks, but only for assigned host groups, clouds and library servers. For example, if you wanted to grant a person administrative privileges over a specific host group, cloud or library server, you would assign this role profile and a scope that defined the objects over which you wanted to grant this permission. Members of this group cannot add XenServer hosts or clusters and cannot add WSUS servers to VMM. Role members are unable to modify membership of the Administrator user role. In previous versions of VMM, this role was termed Delegated Administrator.
- **Read-Only Administrator.** Members of this role are able to view settings, status and job status of objects that exist within specific host groups, clouds and library servers. Importantly though, members of this role are unable to change what they can view. You would delegate this role to people who are responsible for auditing configurations.

- **Tenant Administrator.** Members of this role are able to manage self-service users and VM networks. A user delegated the Tenant Administrator role profile is able to create, deploy and manage their own VMs. They can also specify the tasks that self-service users are able to perform on self-service user-owned VMs and services. A tenant administrator can configure quotas on resources such as RAM, CPU and network as well as VMs within the scope they have been assigned.
- **Application Administrator (Self-Service User).** Members of this role are able to create, deploy and configure their own VMs within a specific scope. When configuring a user role that includes the application administrator role profile, you can specify quotas, such as limiting the number of VMs of a specific type that the user who is assigned the role can deploy. In prior versions of VMM, this role profile was named Self-Service User.

Role members

You can configure roles by selecting individual user accounts or Active Directory security groups.

Individual user accounts, either directly or through Active Directory security group membership, can be members of multiple roles. When assigned multiple roles, delegated privileges are cumulative. Good practice when assigning role members is to assign to Active Directory security groups rather than to individual user accounts. This way you can add and remove individual users from a user role by modifying the Active Directory security group membership. When creating Active Directory security groups, remember to give the group a name related to the permissions that members of the group will have with the user role. Figure 5b shows addition of the Veeam.VMM.Sydney.Admins Active Directory security group to a new VMM user role.

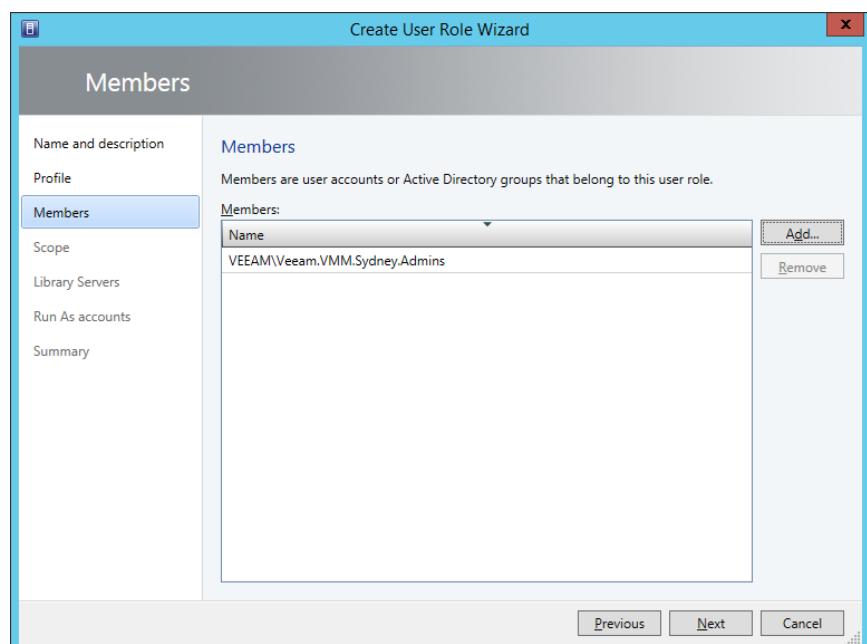
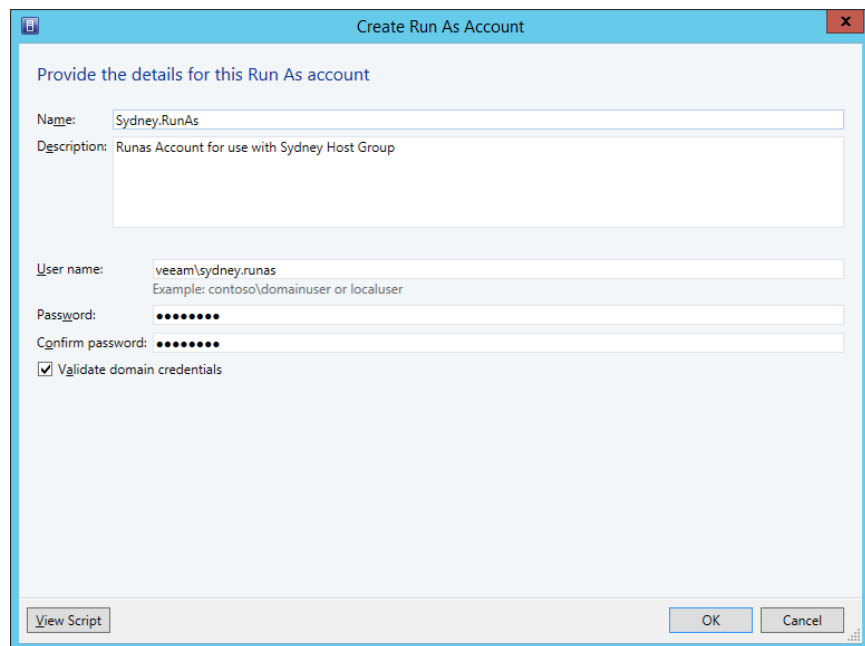


Figure 5b. Role members

Run As accounts

Run As accounts are pre-configured accounts, including username and password, that can be used to perform specific VMM tasks. Usually when a user attempts to perform a task in VMM, the permissions of the currently signed-on user will be used. When a user performs tasks in VMM with a Run As account, the permissions of the Run As account will be used. Figure 5c shows the creation of a Run As account named Sydney.RunAs, which maps to the user account `veeam\sydney.runas`.



The screenshot shows a Windows-style dialog box titled "Create Run As Account". The dialog has a light blue header and a white body. The text "Provide the details for this Run As account" is displayed at the top. Below this, there are several input fields: "Name" with the value "Sydney.RunAs", "Description" with the value "Runas Account for use with Sydney Host Group", "User name" with the value "veeam\sydney.runas" and a small example text "Example: contoso\domainuser or localuser" below it, "Password" and "Confirm password" both masked with black dots. A checkbox labeled "Validate domain credentials" is checked. At the bottom of the dialog, there are three buttons: "View Script", "OK", and "Cancel".

Figure 5c. Run As account

Scope

Scopes allow you to define the objects with which the permissions you are delegating can be used. If configuring a user role using the Fabric Administrator or Read-Only Administrator role profile, you can set scopes using existing private clouds or host groups. If you are configuring a user role using the Tenant Administrator or Application Administrator role profiles, you'll only be able to select private clouds. Figure 5d shows scope selection when creating a user role.

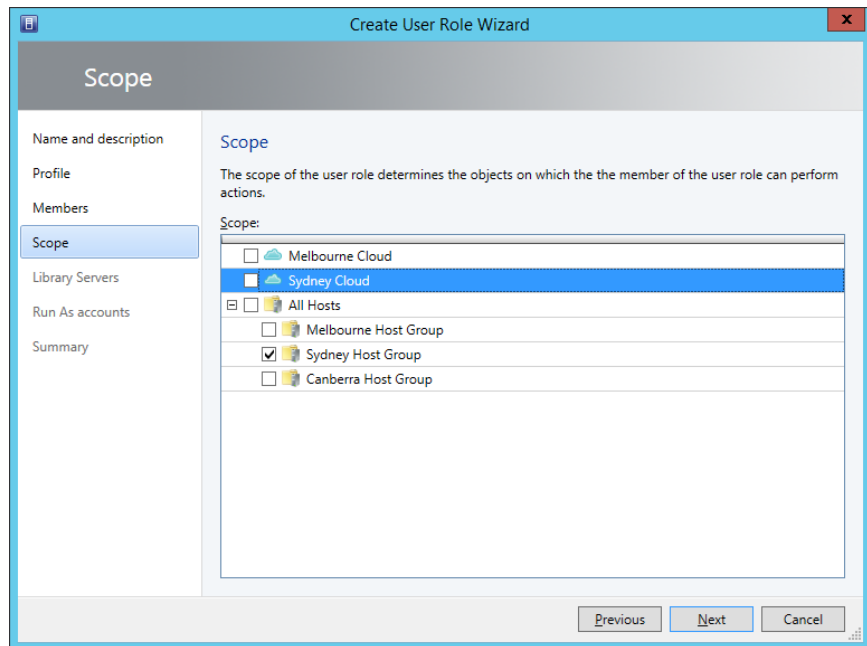


Figure 5d. Configure scope

Quotas

Quotas allow you to limit the use of virtual CPUs, memory, storage and VMs for user roles that use the Tenant Administrator and Application Administrator profiles. You can configure quotas at the role level and at the member level. When configured at the role level, all members of the role can use resources up to the specified limit. When configured at the member level, the limit applies to each role member. Figure 5e shows quota configuration when creating a user role.

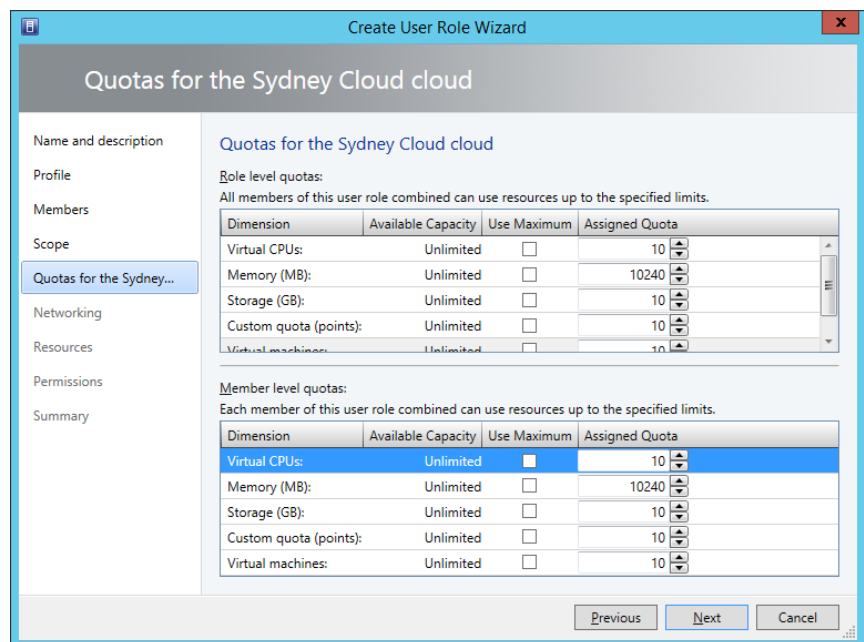


Figure 5e. Configure quotas

Networks and resources

When configuring the Tenant Administrator and Application Administrator profiles, you can specify which VM networks and resources a role member may use. Resources include templates and profiles. For example, if you want to limit members of a user role so that they can only use a specific guest OS profile and template, you would configure those when configuring resources when setting up the user role. Figure 5f shows available VM templates, hardware profiles and guest OS profiles.

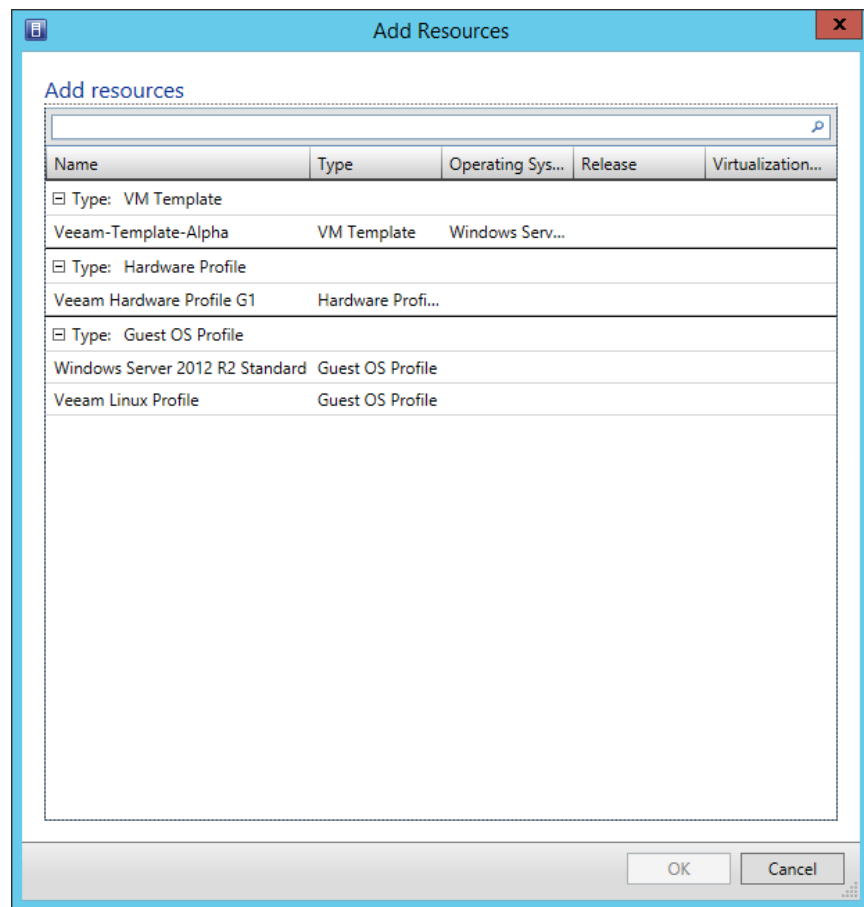


Figure 5f. Specify resources

Permissions

When configuring the Tenant Administrator and Application Administrator profiles, you can specify the permissions that the person assigned the user role has with regard to the objects within the scope of the user role. For example, you can assign the Shut Down permission, which would allow the person assigned the user role to shut down VMs within the private cloud associated with the user role. Figure 5g shows some of the permissions you can assign when creating a user role.

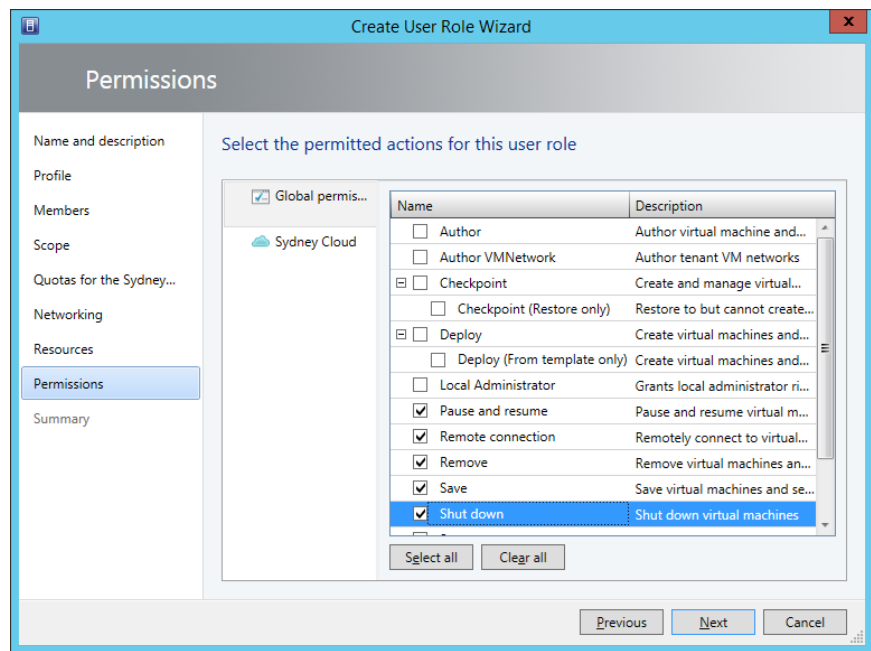


Figure 5g. User role permissions

The available permissions are as follows:

- **Author.** Allows user to create VM and service templates.
- **Author VMNetwork.** Allows user to author tenant VM networks.
- **Checkpoint.** Allows user to create and manage VM checkpoints.
- **Checkpoint (Restore only).** Allows user to restore a VM from an existing checkpoint, but does not allow the creation of a new checkpoint.
- **Deploy.** Allows creation of VMs and services from virtual hard disks or templates
- **Deploy (from template only).** Allows creation of VMs, but only from templates.
- **Local Administrator.** Grants local Administrator rights within the VM OS.
- **Pause and resume.** Allows user to pause and resume VMs.
- **Remote connection.** Allows user to remotely connect to VMs.
- **Remove.** Allows user to delete VMs and services
- **Save.** Allows user to save VMs and services
- **Shut down.** Allows user to shut down VMs.
- **Start.** Allows user to start VMs.
- **Stop.** Allows user to stop VMs.
- **Store and re-deploy.** Allows user to store VMs in the library and re-deploy VMs that the user has stored.

Find Out More:

You can learn more about User Roles in Virtual Machine Manager by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg696971.aspx>

Managing VMM libraries

VMM libraries are a collection of file-based and non-file-based resources. File-based resources in a library can include virtual hard disks, ISO images, scripts, driver files and application packages (SQL Server data-tier applications, Web Deploy Packages and Server App-V packages). Non-file-based library resources include VM templates, service templates, and hardware and guest OS profiles. These non-file-based library resources are stored in the VMM database.

To use a file-based resource with VMM, you need to add that resource to a VMM library. For example, to use a specific specially prepared virtual hard disk with VMM deployments, you first add that virtual hard disk to a VMM library. When you add the file to the VMM library, an automatic discovery process runs, discovering the resource and adding it to the library. The content in VMM libraries does not replicate automatically, so it may be necessary to add the same file-based resource to multiple libraries to ensure that it is available in multiple locations.

It is important to note that library servers are only able to discover files associated with versions of the OS equal to or earlier than the library server OS. For example, only library servers running the Windows Server 2012 or Windows Server 2012 R2 can discover virtual hard disk files in .vhdx format. Library servers running Windows Server 2008 R2 will not recognize these files as virtual hard disks.

When you deploy VMM, you can create a library share or select an existing library share. By default, the setup will configure a library share named *MSSCVMMLibrary* on the VMM management server. Library servers have the following prerequisites:

- Must be in the same domain, or in a domain that has a two-way trust relationship with the domain that hosts the VMM management server.
- The firewall must be configured to allow File and Print Sharing (SMB) traffic.
- You must create a shared folder on the server that will host the library share before you configure the server to be a VMM library server or host a VMM library share.
- The account used to add a library server must have local administrator rights on the computer that will host the library server.

Equivalent objects

Equivalent objects allow you to mark specific file-based resources stored in VMM libraries as equivalent. For example, you could have a virtual hard disk file that hosts a sysprepped deployment of Windows Server 2012 R2 stored in library shares at your organization's Sydney and Melbourne sites. If you mark these virtual hard disk files as equivalent, you can then create a new VM template that uses this virtual hard disk file, allowing you to use the same template across multiple locations. Figure 5h shows two hard disks stored on different library servers that are marked as equivalent.

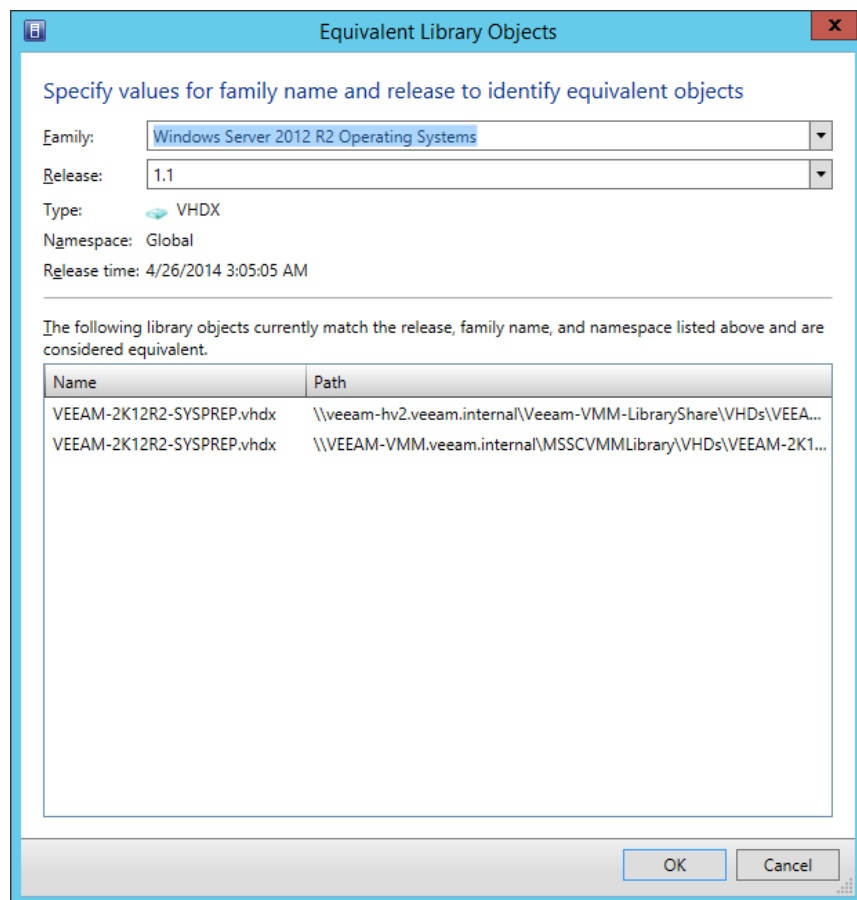


Figure 5h. Equivalent objects

Host group libraries

You can associate VMM libraries with specific host groups. VMM will use this information when determining which resources to use if you have specified a resource with an equivalent object in a profile or template. For example, a virtual hard disk file marked as equivalent and hosted on shares on the Melbourne and Sydney VMM library servers is specified as a resource in a template. If you associate the Sydney VMM library server with the Sydney host group, future deployments to the Sydney host group using that template will use template resources that are equivalent stored on the Sydney VMM library server. Figure 5i shows the library server VEEAM-VMM.veeam.internal associated with the Melbourne host group.




Figure 5i. Library server association

Find Out More:

You can learn more about VMM Libraries by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610598.aspx>

Integrate third-party virtualization platforms

You can use System Center 2012 R2 VMM to manage VMware ESX, VMware ESXi and Citrix XenServer virtualization hosts. To manage VMware virtualization hosts, you must first have deployed a VMware vCenter server. You can then configure VMM to connect to the vCenter server and, through that server, to perform management tasks on VMware ESX and VMware ESXi virtualization hosts.

System Center 2012 R2 VMM supports the following versions of VMware vCenter:

- VMware vCenter Server 4.1
- VMware vCenter Server 5.0
- VMware vCenter Server 5.1

When connected through a supported version of VMware vCenter, System Center 2012 R2 VMM can manage VM hosts and host clusters running the following versions of VMware ESX and ESXi

- VMware ESXi 5.0
- VMware ESXi 5.1
- VMware ESXi 4.1
- VMware ESX 4.1

To manage Citrix XenServer hosts, you need to add the Citrix XenServer – Microsoft System Center Integration Pack. When you do this, you can use System Center 2012 R2 VMM to manage versions 6.0 and 6.1 of Citrix XenServer.

Find Out More:

You can learn more about using VMware ESX and Citrix XenServer with VMM by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610687.aspx>

Bare metal Hyper-V host deployment

You can use VMM to perform bare metal Hyper-V host deployment. VMM can perform this task using the PXE Server component of the Windows Deployment Services role, which you can deploy on computers running Windows Server operating systems. You do this by configuring physical computer profiles (previously termed host profiles).

Before you can create a physical computer profile, you need the following:

- Sysprepped/generalized virtual hard disk deployed to an existing library share
 - In .vhd format, this supports Windows Server 2012, Windows Server 2008 R2 with SP1 and Windows Server 2008.
 - In .vhdx format, support for Windows Server 2012 and Windows Server 2012 R2
- Any drivers required to support the hardware that the Hyper-V host will be deployed to must be present in the VMM library.
- If you want the Hyper-V host to use a static IP address assigned through VMM, you must ensure that the logical network that the host will use has a network site associated with it. This network site will need a static IP address pool managed by VMM.

- Any answer files or scripts required with Hyper-V host deployment must be present in the VMM library.
- You must configure a Run As account that has been delegated necessary privileges to join the Hyper-V host to a domain that is trusted by the domain that hosts the VMM management server.
- Ensure that the bare metal hardware chassis that will be the target of the deployment is configured to PXE boot and has a processor that is configured to support Hyper-V virtualization.
- You must have deployed a PXE server and added it to VMM on the subnet that hosts the bare metal hardware chassis that you want to configure as a Hyper-V virtualization host.

Find Out More:

You can learn more about Physical Computer Profiles by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610653.aspx>

Integrating Operations Manager with VMM

You can use System Center 2012 R2 Operations Manager to monitor your organization's virtualization infrastructure. Connecting System Center 2012 R2 Operations Manager with System Center 2012 R2 VMM provides you with the ability to:

- Monitor the health and availability of:
 - The VMM management server
 - The VMM database server
 - VMM library servers
 - Virtualization hosts
- See diagram views of the virtualized environment
- Use Performance and Resource Optimization (PRO)
- Use maintenance mode integration
- Use SQL Server Analysis Services for VMM

To connect Operations Manager with VMM, you need to:

- Install the Operations Manager console on the VMM management server
- Import the following management packs into Operations Manager:
 - Windows Server Internet Information Services 2003
 - Windows Server 2008 Internet Information Services 7
 - Windows Server Internet Information Services Library
 - SQL Server Core Library
- Deploy Operations Manager agent on the VMM management server and all virtualization hosts managed by the VMM server. You must ensure that virtualization hosts are configured as shown in Figure 5j, by selecting the checkbox: **Allow this agent to act as a Proxy and discover managed objects on other computers.**

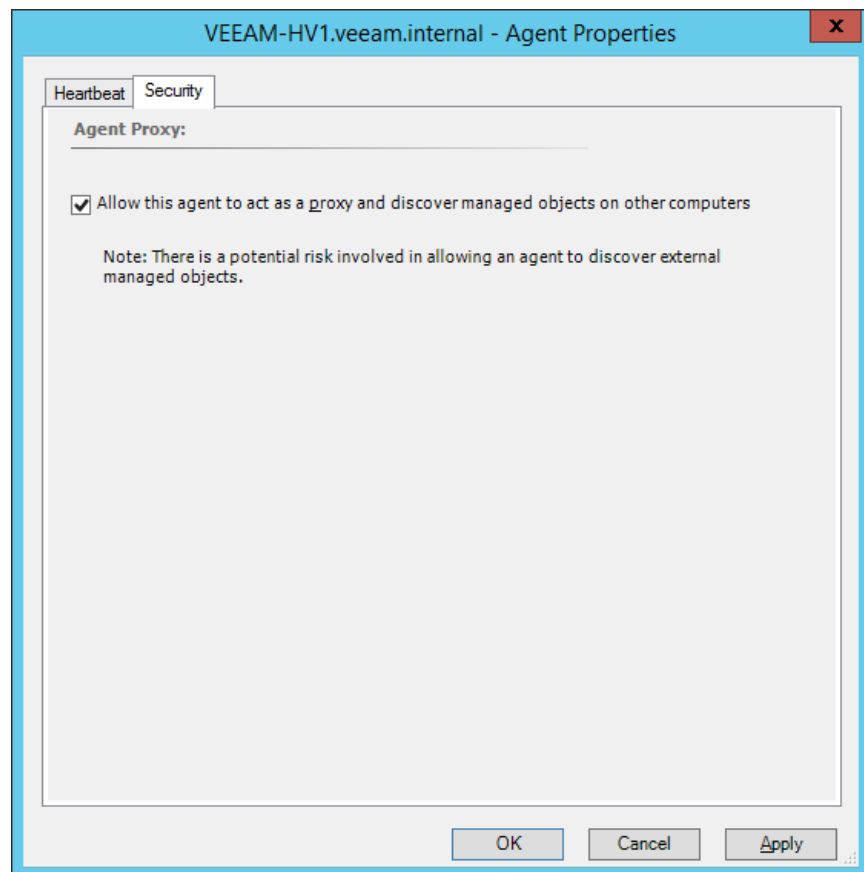


Figure 5j. Configure agent to act as proxy

Once you have configured the prerequisites, perform the following steps to configure integration:

1. In the Settings workspace of the VMM console, right-click Operations Manager server and click Properties.
2. On the Connection to page, specify the server name of the Operations Manager server, an account with administrative privileges for the management group you want to allow VMM to connect to, and whether you want to enable PRO and maintenance mode integration. This page is shown in Figure 5k.

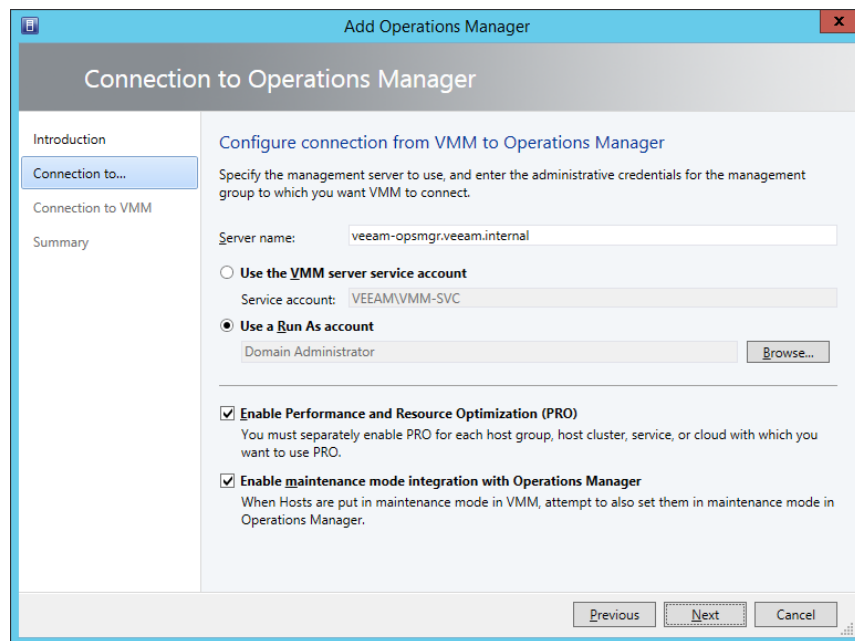


Figure 5k. Configure connection to Operations Manager

3. On the Connection to VMM page, specify the credentials of the account used by Operations Manager to connect to VMM.
4. Complete the wizard. Once the wizard is complete, you can view the properties of the Operations Manager server and verify that the connection status is set to **OK** as shown in Figure 5l.

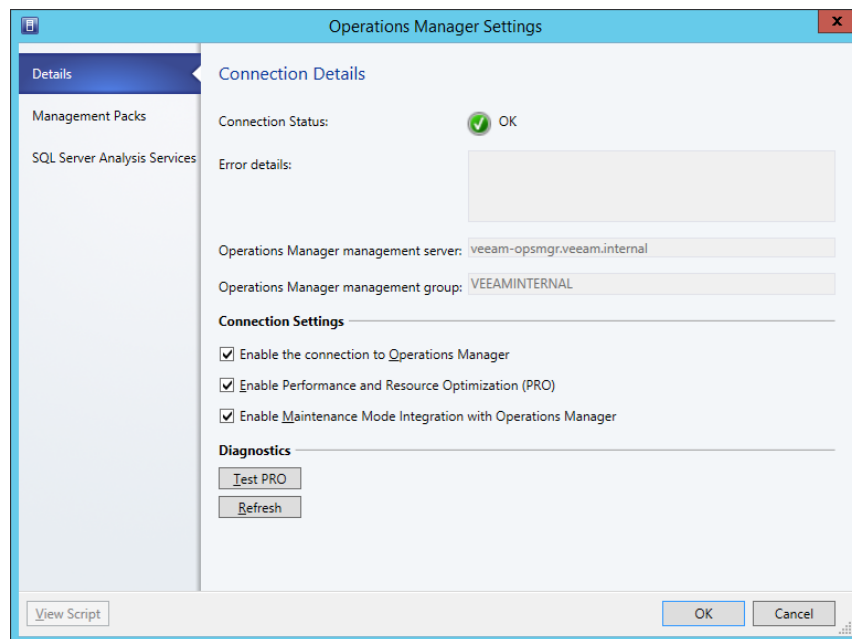


Figure 5l. Verify Operations Manager connection

- When connected, the monitoring workspace of the Operations Manager console will contain the VMM, VMM PRO and VMM Views items as shown in Figure 5m.

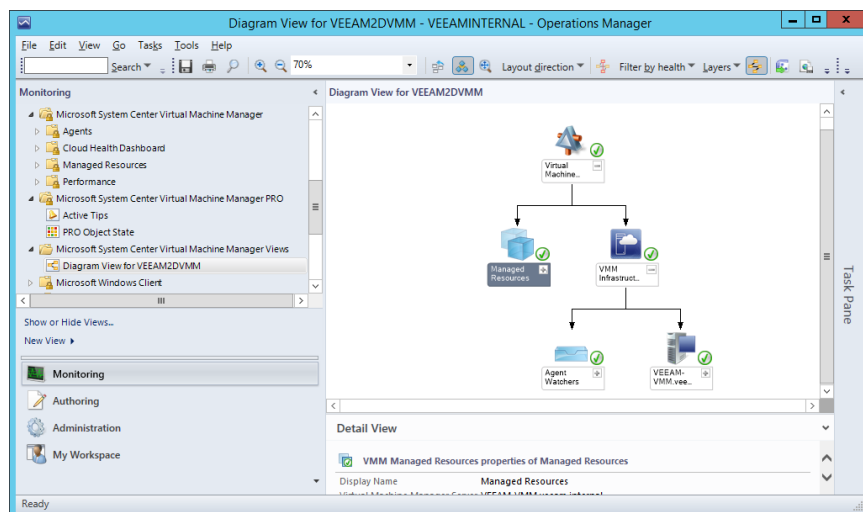


Figure 5m. Virtual Machine Manager views

Find Out More:

You can learn more about configuring Operations Manager integration with VMM by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh427287.aspx>

Integrating Service Manager with VMM

System Center 2012 R2 Service Manager allows you to implement service management as defined by the Information Technology Infrastructure Library (ITIL) and Microsoft Operations Framework (MOF). By deploying Service Manager, you can have a central platform to manage incidents, problems, changes and releases, not just for your organization's VMM infrastructure, but for your organization's entire IT infrastructure.

There are two ways of importing data about your organization's VMM infrastructure into Service Manager. The first is to configure the Operations Manager connector for Service Manager. You use this when you have integrated VMM with Operations Manager as described earlier in this chapter. While this will import the majority of available information about your organization's VMM infrastructure into Service Manager, for complete coverage, you should configure the VMM connector for Service Manager. This connector allows you to import objects, including VM templates, service templates, and storage classifications into Service Manager. To configure the VMM connector for Service Manager, perform the following steps:

1. In the Administration workspace, click Connectors.
2. In the Tasks pane, click Create Connector and click Virtual Machine Manager Connector.
3. On the General page of the VMM connector wizard, specify a name for the connector.
4. On the Connection page, enter the server name and a Run As account that has permissions to establish the connection, as shown in Figure 5n. You can use the Test Connection to verify that the connection works before creating it.

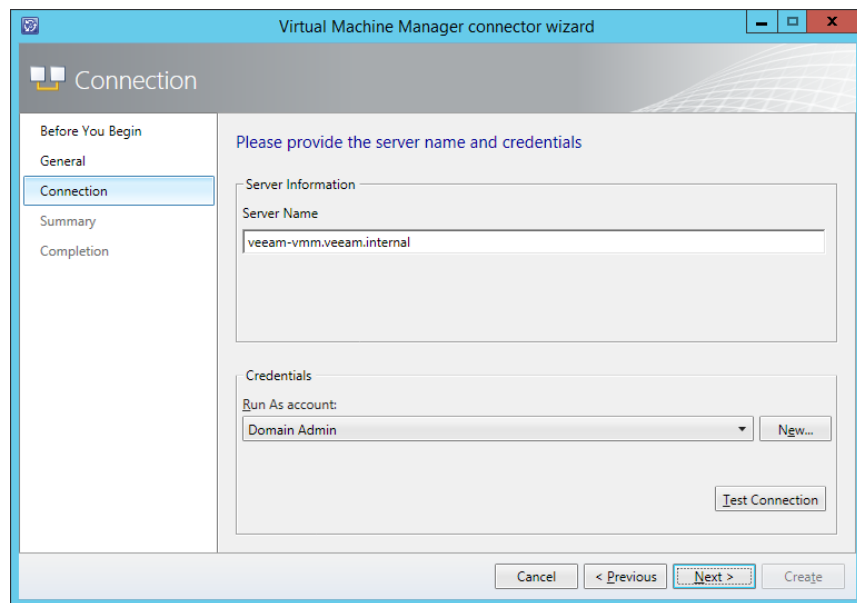


Figure 5n. Service Manager connection to VMM

Find Out More:

You can learn more about this topic by consulting the following TechNet Wiki article: http://social.technet.microsoft.com/wiki/contents/articles/12323.system-center-2012-integration-guide-virtual-machine-manager.aspx#Service_Manager

Servicing virtual machine images

Rather than deploying VMs and then having those VMs contact a software update server to acquire recently released software updates, you can instead apply the updates directly to VM images stored in the VMM library. When you do this, VMs deployed from these images will be up to date at the time of deployment.

There are several tools that you can use to perform this task:

- Dism.exe
- Virtual Machine Servicing Tool
- Orchestrated Offline VM Patching Runbook

Dism and Add-WindowsPackage

You can use the `dism.exe` utility or the **Add-WindowsPackage** Windows PowerShell cmdlet to inject updates into a VM image in `.vhd` and `.vhdx` format. To use this method, you must download the update files required for each image from Microsoft's website, mount each image and then apply the updates to each image using a command similar to:

Dism.exe /image:c:\mount /Add-Package /PackagePath:"c:\updates\Security Update for Windows Server 2012 R2 (KB2893294)"

This method is labor intensive, requires you to manually apply updates to each image, and is unsuitable for organizations that have a large number of operating system images stored within the VMM library.

Virtual Machine Servicing Tool 2012

The Virtual Machine Servicing Tool (VMST) 2012 works with System Center 2012 VMM, but not officially with System Center 2012 R2 VMM. You can use VMST 2012 to apply updates to:

- Offline VMs in a SCVMM library
- VM templates
- Offline virtual hard disks
- Stopped and saved state VMs on a virtualization host

To use VMST 2012, you need to perform the following steps:

- Deploy the tool. During deployment, you will connect the tool to your VMM management server and your software update server (WSUS or System Center 2012 Configuration Manager).
- Specify the VMs, templates and virtual hard disk groups the VMST will update.
- Create and schedule servicing jobs. These jobs allow you to specify which VMs, templates and virtual hard disk groups should be updated, what resources to use for the updates and when to perform the updates.

The drawback with VMST 2012 is that there is no official support for using the product with System Center 2012 R2 VMM, with only support for System Center 2012 VMM listed on the tool's TechNet Page.

Find Out More:

You can learn more about automatically patching images in the Virtual Machine Manager library using the VMST by consulting the following TechNet article:

<http://technet.microsoft.com/en-us/library/jj149757.aspx>

Orchestrated Offline VM Patching Runbook

The Orchestrated Offline VM Patching Runbook is a tool made available through TechNet that you can use to automate the updating of operating system images on virtual hard disks stored on VMM library servers. When deployed, this runbook will update images in the VMM library as soon as you approve them for deployment through your organization's WSUS-based software update solution. The Orchestrated Offline VM Patching Runbook requires Windows Server 2012 R2, System Center 2012 R2 VMM and the Windows Azure Pack. To use this solution, you need to have installed and registered Service Management Automation in the Windows Azure Pack Admin Portal.

The virtual hard disks that host the operating system to be patched must be running Windows Server 2008 SP1, Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2 and the virtualization platform must be set in VMM to "Microsoft Hyper-V."

This method requires that you deploy and configure Windows Azure Pack alongside your existing VMM deployment.

Find Out More:

You can learn more about automatically patching images in the VMM library by consulting the following TechNet blog post: <http://blogs.technet.com/b/privatecloud/archive/2013/12/07/orchestrated-vm-patching.aspx>

Integrating Data Protection Manager

System Center 2012 R2 Data Protection Manager (DPM) is the System Center 2012 R2 product that allows you to perform backup and recovery tasks. You can use DPM to perform backup and recovery of multiple Hyper-V virtualization hosts as well as configure protection for applications and data hosted within VMs. DPM supports backup replication to a second DPM server. DPM uses disk to store protected short-term data. You can configure DPM to write protected data to tape for long-term protection. You can also configure DPM to replicate protected data to Windows Azure, where it can be stored for up to 120 days.

For example, to configure DPM to protect the VMs hosted on one or more Hyper-V virtualization hosts, you first deploy the DPM agent to each server running Hyper-V. You then create a DPM protection group, specify which VMs you want to back up, as shown in Figure 5o, and then configure protection settings (such as how long you want to store protected data, how frequently DPM performs synchronization and how often DPM performs consistency checks on protected data).

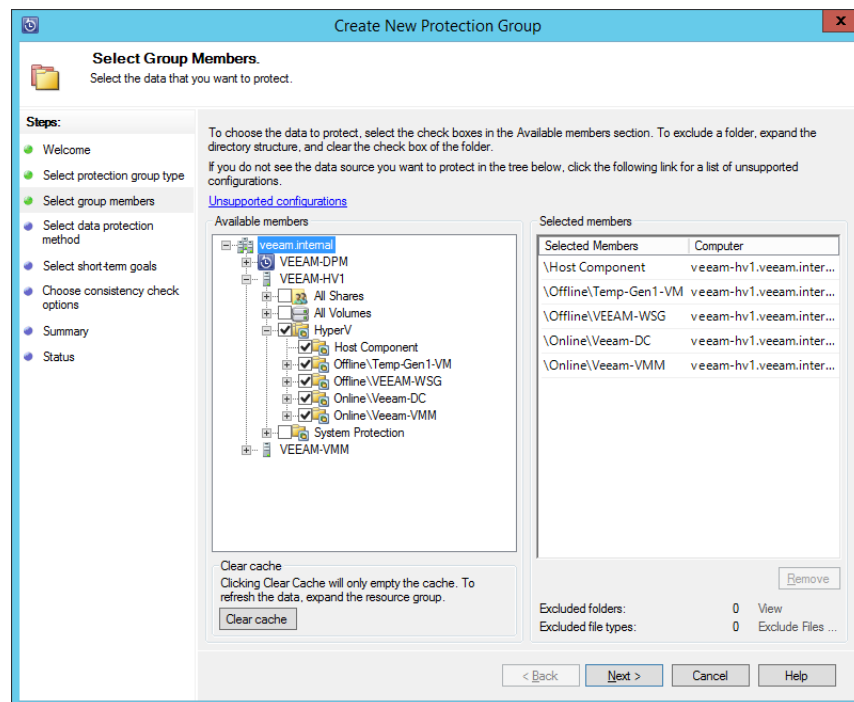


Figure 5o. Protect VMs using DPM

Once you've configured VMs on a Hyper-V host to be protected, you can also configure a script that runs as a scheduled task, which automatically adds any new VMs deployed to a protected Hyper-V server to a DPM Protection Group.

You can deploy DPM agents to individual VMs, allowing protection within the VM. For example, in heterogeneous environments where you might be using virtualization hosts running Hyper-V and a third-party solution, you may choose to back up Windows VMs running on the third-party virtualization hosts by deploying the agent to the VM.

To protect a VMM management server, you need to protect the VMM database as well as any library shares that the server hosts. Figure 5p shows the configuration of a DPM Protection Group for a VMM management server where the VMM database and the library share are configured for protection.

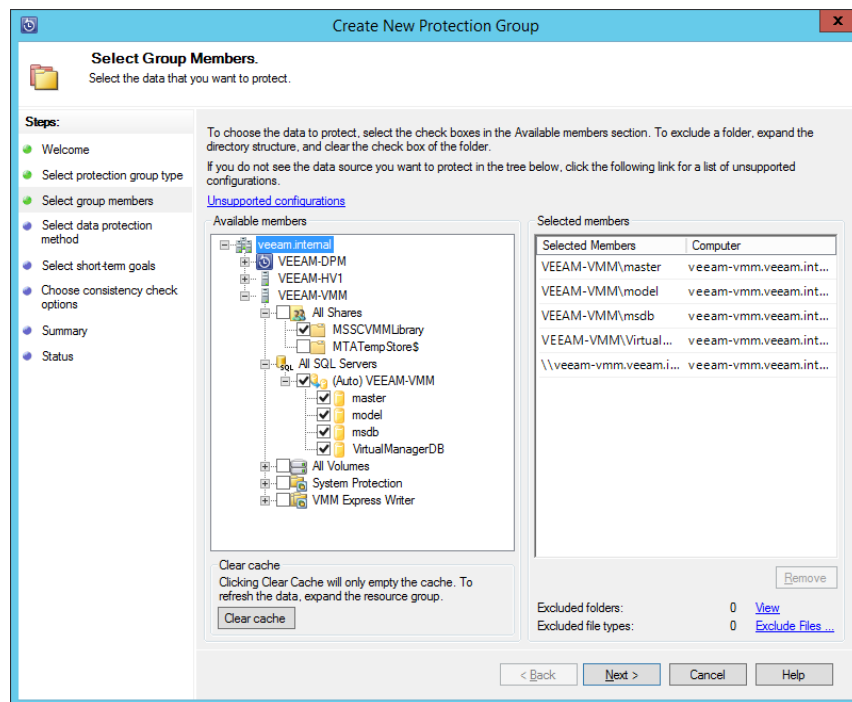


Figure 5p. Protect VMM database and library

Find Out More:

You can learn more about protecting VMM hosts with Data Protection Manager by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh771597.aspx>

DPM supports the following VM recovery options as shown in Figure 5q:

- Recovery of the VM to the original location. When you select this option, the original VM's hard disk is deleted. DPM will recover both the virtual hard disks and VM configuration files to their original location.
- Recovery of the VM to an alternate location. When you select this option, you can recover the VM to a different Hyper-V host that has the DPM agent installed.
- Recovery to a network folder. When you choose this option, you can recover the VM files to a network folder.

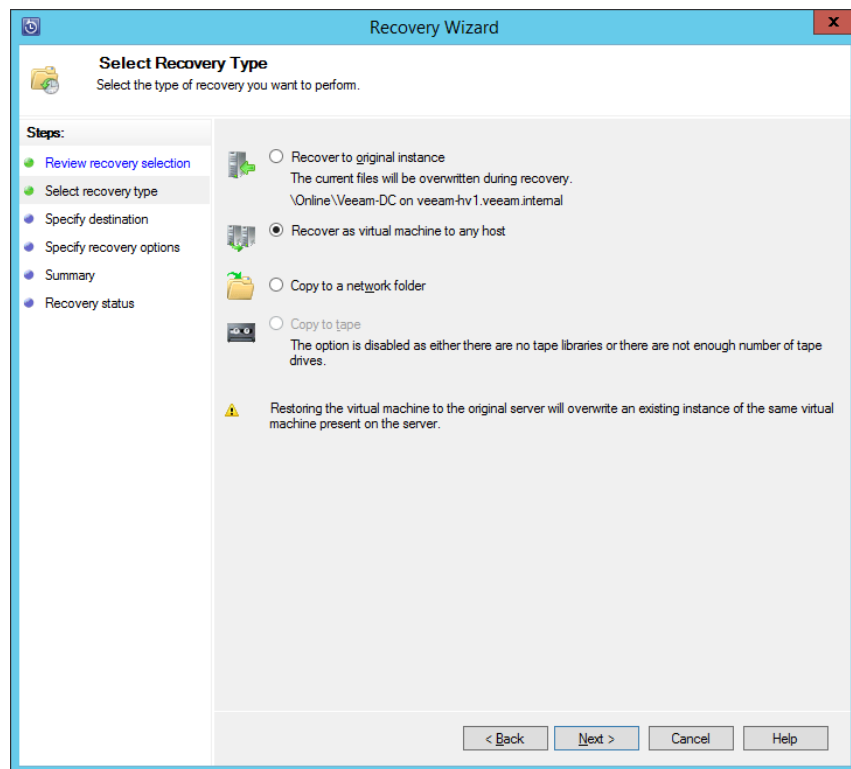


Figure 5q. VM recovery location

DPM also supports item-level recovery. Item-level recovery allows you to recover specific files, folders, volumes and virtual hard disks from a protected VM to a network share or volume on a server that has the DPM agent installed. When you perform item-level recovery, you are only able to recover files, folders, volumes and virtual hard disks to a network location.

Find Out More:

You can learn more about recovering VMs with Data Protection Manager by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh757981.aspx>

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. You want to create a user role in VMM that allows users that are members of the role the ability to perform all administrative tasks, but only over the Sydney and Melbourne host groups. Which role profile would you select when creating the user role?
2. You want to create a user role in VMM that allows users that are members of the role to view all settings, status and jobs status information related to VMM objects. Users that are members of these roles should not be able to configure settings. Which role profile would you select when creating the user role?
3. You are creating a user role in VMM. Which role profiles can only be assigned private cloud scopes and cannot be assigned host group scopes?
4. You want to ensure that a single member of a user role can only use a maximum of 10 virtual CPUs. What type of quota should you configure when configuring the user role?
5. You want to configure a quota for the Sydney Development Users user role so that all role members combined cannot consume more than 5 TB of storage. What type of quota should you configure when configuring the user role?
6. You have an identical sysprepped/generalized virtual hard disk deployed to VMM library servers at five different branch offices. You want to configure a single VM template in such a way that the copy of the virtual hard disk hosted on the local branch office VMM library server is used when a VM is deployed from the template. What steps must you take to accomplish this goal?
7. On which server must you install the Operations Manager console if you want to integrate Operations Manager with VMM?

Chapter 6: Hyper-V Failover Clustering and Failover Clustering Roles

Windows Server 2012 and Windows Server 2012 R2 support Hyper-V failover clusters with up to 64 nodes hosting up to 8,000 virtual machines (VMs). Before you deploy Windows Server 2012 R2 Hyper-V failover clusters, you should have a good working knowledge of topics such as shared storage, quorum, networking and cluster upgrades. You should also know about technologies new to Windows Server 2012 and Windows Server 2012 R2 such as Cluster Aware Updating, Cluster Shared Volumes, Active Directory detached clusters, and shared virtual hard disks.

In this chapter you'll learn about:

- Cluster shared storage
- Cluster quorum
- Cluster networking
- Force Quorum Resiliency
- Cluster Aware Updating
- Upgrading clusters
- Cluster Shared Volumes
- Active Directory detached clusters
- Scale-out file servers
- Preferred owner and failover settings
- Guest clustering
- Shared virtual hard disk

Cluster shared storage

Windows Server 2012 R2 Hyper-V failover clusters require some form of shared storage to store VM hard disk and configuration files. Windows Server 2012 R2 Hyper-V failover clusters can use SAS, iSCSI, Fibre Channel or Fibre Channel over Ethernet (FCoE) for shared storage.

Consider the following when choosing storage for a Hyper-V failover cluster:

- SAS is suitable for two-node failover clusters where the cluster nodes are physically close together (less than 10 meters apart). If your failover cluster nodes are further apart or there are more than two, consider another option.
- iSCSI works over Ethernet and works with the maximum of 64 nodes supported by a Windows Server 2012 R2 Hyper-V failover cluster. Windows Server 2012 and Windows Server 2012 R2 have an iSCSI Target Server feature that allows servers running the OS to function as iSCSI storage .
- Fibre Channel requires that each cluster node be configured with a special Host Bus Adapter to connect to shared storage. This often makes Fibre Channel a more expensive option than iSCSI.
- FCoE encapsulates Fibre Channel traffic over an Ethernet network. FCoE can be used with network adapters that include the Host Bus Adapter and NIC functionality, or with special network adapters that allow FCoE encapsulation to be performed in software.

Failover cluster storage has the following conditions:

- Use of dynamic disks is not supported for Hyper-V failover cluster shared storage. Shared storage should use basic disks.
- Disk witnesses can be formatted as NTFS or ReFS volumes.
- Ensure that storage is isolated so that separate failover clusters are unable to access the same LUN. You can accomplish this with LUN masking or zoning.
- Where possible, use Multipath I/O (MPIO) and teamed network adapters to ensure that there are multiple paths to storage.

Find Out More:

You can learn more about failover cluster storage options by consulting the following TechNet article: <http://technet.microsoft.com/library/jj612869.aspx>

Cluster quorum

Failover clusters remain functional until they do not have enough active votes to retain quorum. Votes can consist of nodes that participate in the cluster as well as disk or file share witnesses. The calculation on whether the cluster maintains quorum is dependent on the cluster quorum mode. When you deploy a Windows Server failover cluster, one of the following modes will automatically be selected, depending on the current cluster configuration:

- Node Majority
- Node and Disk Majority
- Node and File Share Majority
- No Majority: Disk Only

You can change the cluster mode manually or, with Dynamic Quorum in Windows Server 2012 R2, the cluster mode will change automatically when you add or remove nodes, a witness disk or a witness share.

Node Majority

The Node Majority cluster quorum mode is chosen automatically during setup if a cluster has an odd number of nodes. When this cluster quorum mode is used, a file share or disk witness is not used. A failover cluster will retain quorum as long as the number of available nodes is more than the number of failed nodes that retain cluster membership. For example, if you deploy a nine-node failover cluster, the cluster will retain quorum as long as five cluster nodes are able to communicate with each other.

Node and Disk Majority

The Node and Disk Majority model is chosen automatically during setup if the cluster has an even number of nodes and shared storage is available to function as a disk witness. In this configuration, cluster nodes and the disk witness each have a vote when calculating quorum. As with the Node Majority model, the cluster will retain quorum as long as the number of votes that remain in communication exceeds the number of votes that cannot be contacted. For example, if you deployed a six-node cluster and a witness disk, there would be a total of seven votes. As long as four of those votes remained in communication with each other, the failover cluster would retain quorum. Figure 6a shows selecting a storage witness.

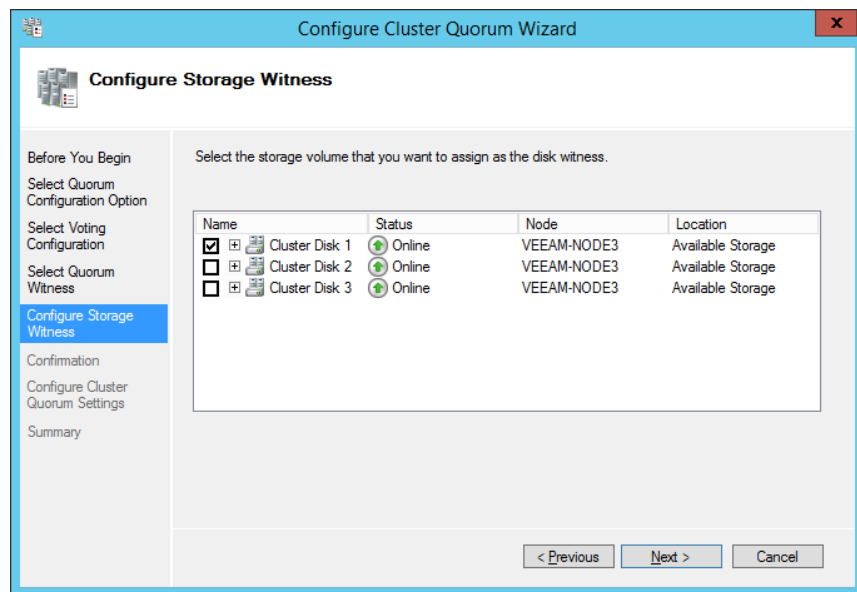


Figure 6a. SELECT DISK WITNESS

Node and File Share Majority

The Node and File Share Majority model is used when a file share is configured as a witness. Each node and the file share have a vote when it comes to determining if quorum is retained. As with other models, a majority of the votes must be present for the cluster to retain quorum. Node and File Share Majority is suitable for organizations that are deploying multi-site clusters; for example, placing half the cluster nodes in one site, half the cluster nodes in another site and the file share witness in a third site. If one site fails, the other site is able to retain communication with the site that hosts the file share witness, in which case quorum is retained.

No Majority: Disk Only

The No Majority: Disk Only model must be configured manually and only used in testing environments as the only vote that counts toward quorum is that of the disk witness on shared storage. The cluster will retain quorum as long as the witness is available, even if every node but one fails. Similarly the cluster will be in a failed state if all the nodes are available, but the shared storage hosting the disk witness goes offline.

Cluster node weight

Rather than every node in the cluster having an equal vote when determining quorum, you can configure which cluster nodes are able to vote to determine quorum by running the Configure Cluster Quorum Wizard. For example, in Figure 6b, cluster node VEEAM-NODE3 is configured so that it does not have a quorum vote. Configuring node weight is useful if you are deploying a multi-site cluster and you want to control which site retains quorum in the event that communication between the sites is lost.

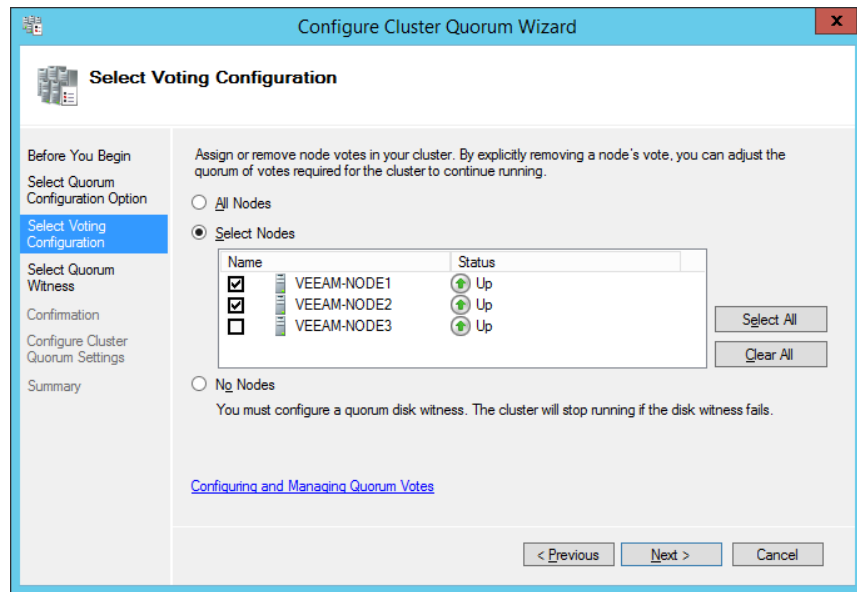


Figure 6b. Configure node votes

You can determine which nodes in a cluster are currently assigned votes by selecting Nodes in the Failover Cluster Manager. Figure 6c shows that VEEAM-NODE3 has an Assigned Vote value of 0

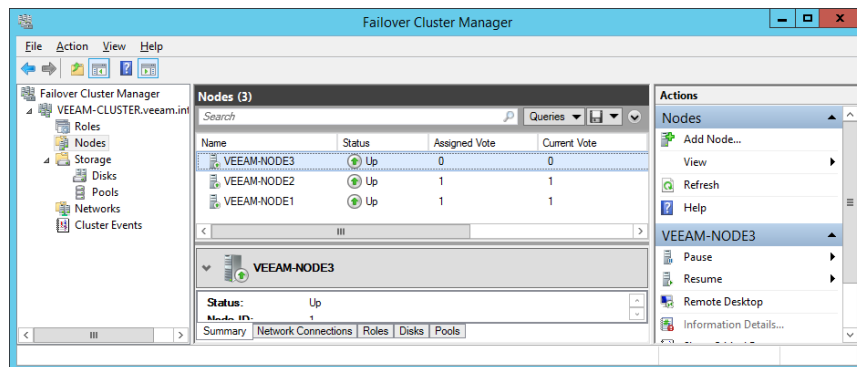


Figure 6c. Verify node votes

Dynamic quorum

Dynamic quorum is a feature introduced with Windows Server 2012 and available in Windows Server 2012 R2 that allows cluster quorum to be recalculated automatically each time a node is removed from or added to a cluster. Dynamic quorum is enabled by default on Windows Server 2012 and Windows Server 2012 R2 clusters. Dynamic quorum works in the following manner:

- The vote of the witness is automatically adjusted based on the number of voting nodes in the cluster. If the cluster has an even number of nodes, the witness has a vote. If a cluster has an even number of nodes and a node is added or removed, the witness loses its vote.

- In the event of a 50% node split, dynamic quorum can adjust the vote of a node. This is useful in avoiding “split brain” syndrome during site splits with multi-site failover clusters.

An advantage of dynamic quorum is that as long as nodes are evicted in a graceful manner, the cluster will reconfigure quorum appropriately. This means that you could change a nine-node cluster so that it was a 5-node cluster by evicting nodes, and the new quorum model would automatically be recalculated assuming that the cluster only had five nodes. With dynamic quorum, it is a good idea to specify a witness even if the initial cluster configuration has an odd number of nodes, because that way a witness vote will automatically be included in the event that an administrator adds or removes a node from the cluster.

Find Out More:

You can learn more about cluster quorum by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj612870.aspx>

Cluster networking

In lab and development environments, it's reasonable to have failover cluster nodes that are configured with a single network adapter. In production environments with mission-critical workloads, you should configure cluster nodes with multiple network adapters, institute adapter teaming and leverage separate networks. Separate networks should include:

- A network dedicated for connecting cluster nodes to shared storage
- A network dedicated for internal cluster communication
- The network that clients use to access services deployed on the cluster

When configuring IPv4 or IPv6 addressing for failover cluster nodes, ensure that addresses are assigned either statically or dynamically to cluster node network adapters. Avoid using a mixture of statically and dynamically assigned addresses as this will cause an error with the cluster validation wizard. Also ensure that cluster network adapters are configured with a default gateway. While the cluster validation wizard will not provide an error if a default gateway is not present for the network adapters of each potential cluster node, you will be unable to create a failover cluster unless a default gateway is present.

Force Quorum Resiliency

Imagine you have a five-node multi-site cluster in Melbourne and Sydney, with three nodes in Sydney. Imagine also that internet connectivity to the Sydney site is lost. Within the Sydney site itself, the cluster will remain running because with three nodes it has retained quorum. But if external connectivity to the Sydney site is not available, you may instead need to forcibly start the cluster in the Melbourne site (which will be in a failed state because only two nodes are present) using the /fq (forced quorum) switch to provide services to clients.

In the past, when connectivity was restored, this would have led to a “split brain” or partitioned cluster, as both sides of the cluster would be configured to be authoritative. To resolve this with failover clusters running Windows Server 2012 or earlier, you would need to manually restart the nodes that were not part of the forced quorum set using the /pq (prevent quorum) switch. Windows Server 2012 R2 provides a feature known as Force Quorum Resiliency that automatically restarts the nodes that were not part of the forced quorum set so that the cluster does not remain in a partitioned state.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK_FQ

Cluster Aware Updating

Cluster Aware Updating (CAU) is a Windows Server 2012 and Windows Server 2012 R2 feature that simplifies the process of applying updates to failover clusters. Prior to the introduction of CAU, it was necessary either to update each node manually or employ a third-party solution to apply software updates to a failover cluster.

CAU uses the following procedure to update a Windows Server failover cluster:

1. CAU has each cluster node obtain update files from the update source. This can be a WSUS server, Configuration Manager or a supported third-party solution.
2. Any workloads hosted on the node are migrated across to other nodes in the cluster and the node is placed in maintenance mode.
3. Software updates are installed on the node. The node restarts if necessary.
4. A check is performed to determine if additional updates are required. If additional updates are found, the updates are installed and the node restarted (if necessary) until the node is up to date.

5. CAU brings the node out of maintenance mode.
6. Workloads hosted on the next node to be updated are migrated to other nodes in the cluster.
7. The next node is placed into maintenance mode, software updates are installed and the process continues until all nodes have been updated.

It is possible to configure CAU so that software updates will be rolled back across all nodes if a single update fails to install on one node in the cluster. This ensures that cluster nodes retain a consistent configuration. You can also configure CAU so that scripts run prior to and after nodes are updated. CAU also supports applying hotfixes and can be customized to apply software updates from vendors other than Microsoft.

You can configure CAU in one of the following modes:

- **Self-updating mode.** When you configure CAU to use self-updating mode, the CAU clustered role is deployed on a cluster node on the cluster that will be updated. During the update process, the CAU role is migrated to another node and continues to manage the process as its original host node is updated. The drawback of this mode is that you are unable to remotely monitor the update process using the CAU interface, although this is possible through the **Get-CauRun** Windows PowerShell cmdlet.
- **Remote-updating mode.** When you configure CAU to use remote-updating mode, the CAU role runs from a computer that is not a member node of the failover cluster. The benefit of this mode is that you can monitor the update process using the CAU interface.

Find Out More:

You can learn more about Cluster Aware Updating by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831694.aspx>

Upgrading clusters

The key to upgrading failover clusters from cluster nodes running Windows Server 2008 and Windows Server 2008 R2 to Windows Server 2012 or Windows Server 2012 R2 is to understand that cluster nodes running Windows Server 2008 and Windows Server 2008 R2 can participate in Windows Server 2012 and Windows Server 2012 R2 failover clusters. However, while cluster nodes running these earlier operating systems can participate in these failover clusters, this configuration is only supported in upgrade scenarios. Let's say, for example, you have a three-node Windows Server 2008 R2 Hyper-V failover cluster that you want to upgrade so that it is a Windows Server 2012 R2 Hyper-V failover cluster. In this scenario you could add three nodes running Windows Server 2012 R2 to the existing cluster, move the VMs from the Windows Server 2008 R2 hosts to the Windows Server 2012 R2 hosts, and then evict the nodes running Windows Server 2008 R2 from the cluster.

In the event that you don't have all that extra hardware, you could instead move all the VMs off one node, evict the node from the cluster, perform an upgrade or a clean install of Windows Server 2012 R2 on the evicted node, and join the node back to the cluster with the new OS. Once the node is returned to the cluster with the new OS, you could transfer part of the cluster workload across to the new node so that another node running the older operating system hosts no workloads, evict that node, upgrade or clean install, rejoin the node, and continue on, completing the migration.

An additional option is to use the Migrate a Cluster Wizard, shown in Figure 6d, to migrate applications or services hosted on a failover cluster running Windows Server 2008 or Windows Server 2008 R2 to a new failover cluster that has nodes running Windows Server 2012 R2.

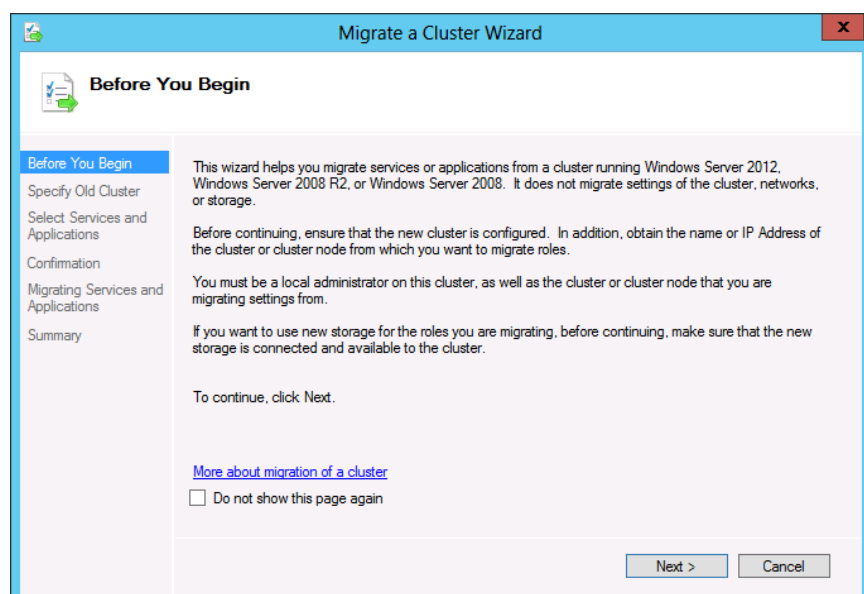


Figure 6d. Migrate a Cluster Wizard

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://blogs.technet.com/b/hugofe/archive/2012/12/06/best-practices-for-migration-of-cluster-windows-2008-r2-2012-as-melhores-praticas-para-migrar-um-cluster-de-windows-2008-para-windows-2012.aspx>

Cluster Shared Volumes

Cluster Shared Volumes (CSVs) are a high-availability storage technology that allows multiple cluster nodes in a failover cluster to have read-write access to the same LUN. This has the following advantages for Hyper-V failover clusters:

- VMs stored on the same LUN can be run on different cluster nodes. This reduces the number of LUNs required to host VMs as the VMs stored on a CSV aren't tied to one specific Hyper-V failover cluster node, but instead can be spread across multiple Hyper-V failover cluster nodes.
- Switch-over between nodes is almost instantaneous in the event of failover because the new host node doesn't have to go through the process of seizing the LUN from the failed node.

CSVs are hosted off Windows Server 2012 and Windows Server 2012 R2 and also allow multiple nodes in a cluster to access the same NTFS or ReFS formatted file system. CSVs support BitLocker, each node performing decryption of encrypted content using cluster computer account. CSVs also integrate with SMB (Server Message Block) Multichannel and SMB Direct, allowing traffic to be sent through multiple networks and to leverage network cards that include Remote Direct Memory Access (RDMA) technology. CSVs can also automatically scan and repair volumes without requiring storage to be taken offline.

You can convert cluster storage to a CSV using the Disks node of the Failover Cluster Manager, as shown in Figure 6e.

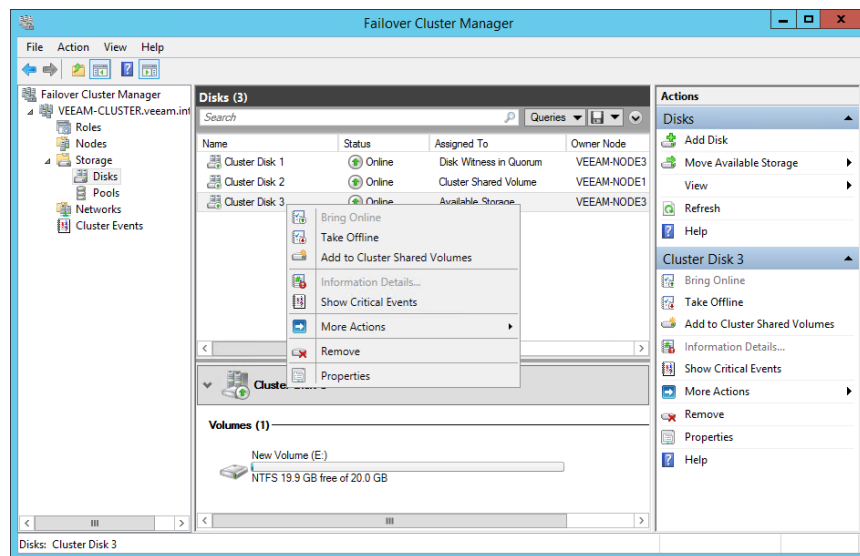


Figure 6e. Add Cluster Shared Volume

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj612868.aspx>

Active Directory detached clusters

Active Directory detached clusters, also termed “clusters without network names,” are a feature of Windows Server 2012 R2. Detached clusters have their names stored in DNS, but do not require the creation of a computer account within Active Directory. The benefit of detached clusters is that it is possible to create them without requiring that a computer object be created in Active Directory to represent the cluster, or that the account used to create the cluster have permissions to create computer objects in Active Directory. Although the account used to create a detached cluster does not need permission to create computer objects in Active Directory, the nodes that will participate in the detached cluster must themselves still be domain joined.

Find Out More:

You can learn more about detached clusters by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn265970.aspx>

Scale-out file servers

Scale-out file servers, also known as continuously available file shares, are a type of file share new to Windows Server 2012 and Windows Server 2012 R2. These file shares are hosted on failover clusters and provide the SMB Transparent Failover feature. The advantage of scale-out file servers is that file server connections are distributed over multiple cluster nodes. For example, rather than all clients attempting to access a specific file share through one node, client access to the file share is balanced across all nodes in the failover cluster. You can host shared virtual hard disks on a scale-out file server share. Figure 6f shows configuring a scale-out file server for application data on a Windows Server failover cluster.

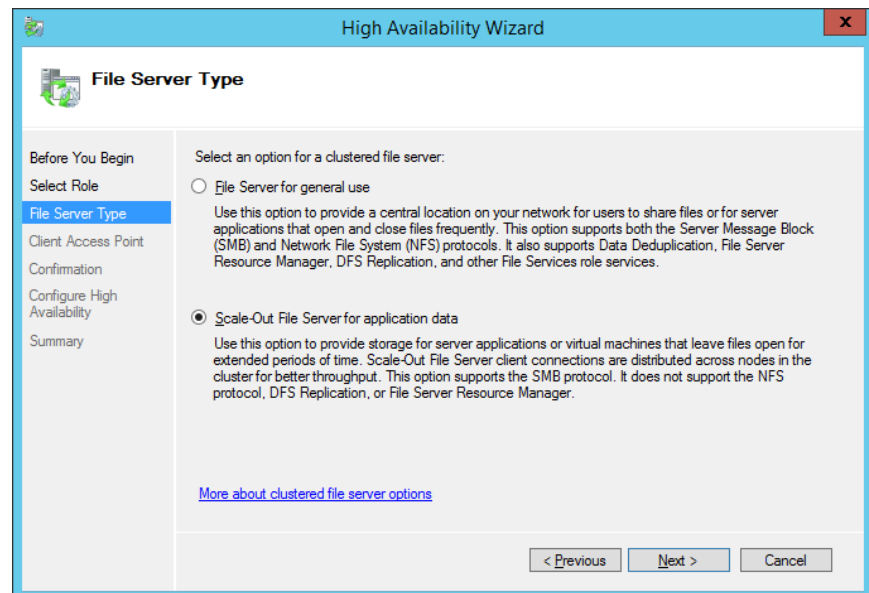


Figure 6f. Configure scale-out file server

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831349.aspx>

Preferred owner and failover settings

Cluster role preferences settings allow you to configure a preferred owner for a cluster role. When you do this, as shown in Figure 6g, the role will be hosted on the node listed as the preferred owner. You can specify multiple preferred owners and configure the order in which a role will attempt to return to a specific cluster node.

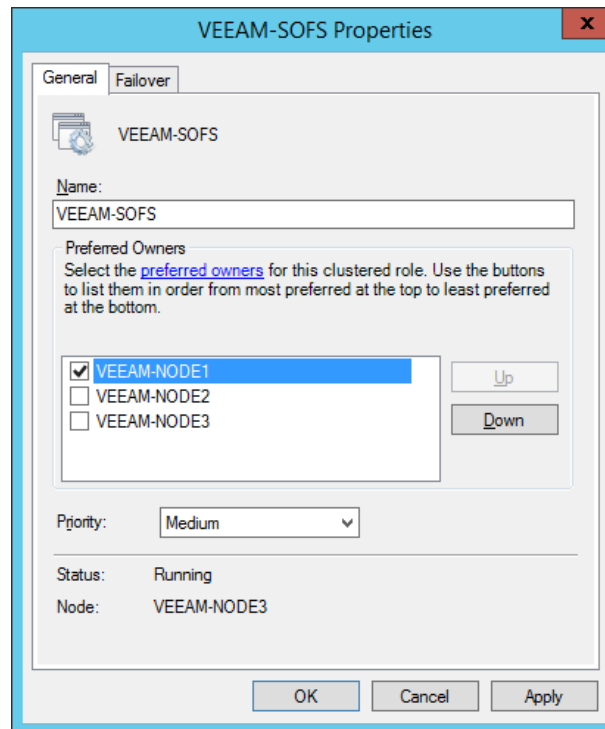


Figure 6g. Configure preferred owner

Failover settings allow you to configure how many times a service will attempt to restart or fail over in a specific period. By default a cluster service can fail over twice in a six-hour period before the failover cluster will leave the cluster role in a failed state. The fallback setting allows you to configure the amount of time a clustered role that has fallen over to a node that is not its preferred owner will wait before falling back to the preferred owner. Figure 6h shows fallback settings.

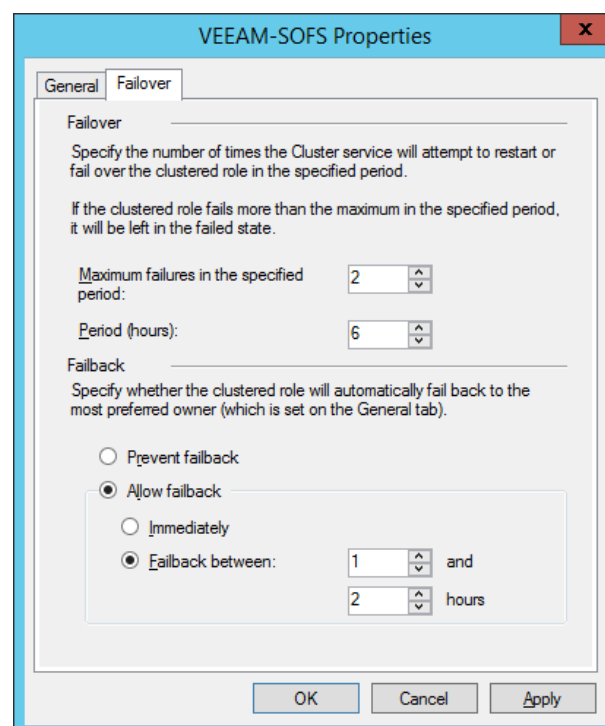


Figure 6h. Fallback settings

Find Out More:

Even though this article refers to Windows Server 2008 R2, you can learn more about preferred owners and failover settings by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/cc771809.aspx>

Guest clustering

A guest cluster is a failover cluster that consists of two or more VMs. You can run a guest cluster on a Hyper-V failover cluster, or you can run guest clusters with nodes on separate Hyper-V failover clusters. While deploying a guest cluster on Hyper-V failover clusters may seem as though it is taking redundancy to an extreme, there are good reasons to deploy guest clusters and Hyper-V failover clusters together, including:

- Failover clusters monitor the health of clustered roles to ensure that they are functioning. This means that a guest failover cluster can detect when the failure of a clustered role occurs and can take steps to recover the role. For example, you deploy a SQL Server 2014 failover cluster as a guest cluster on a Hyper-V failover cluster. One of the SQL Server 2014 instances that participates in the guest cluster suffers a failure. In this scenario, failover occurs within the guest cluster and another instance of SQL Server 2014 hosted on the other guest cluster node continues to service client requests.
- Deploying guest and Hyper-V failover clusters together allows you to move applications to other guest cluster nodes while you are performing servicing tasks. For example, you may need to apply software updates that require a restart to the operating system that hosts a SQL Server 2014 instance. If this SQL Server 2014 instance is participating in a Hyper-V guest cluster, you could move the clustered role to another node, apply software updates to the original node, perform the restart, and then move the clustered SQL Server role back to the original node.
- Deploying guest and Hyper-V failover clusters together allows you to live migrate guest cluster VMs from one host cluster to another host cluster while ensuring clients retain connectivity to clustered applications. For example, a two-node guest cluster hosting SQL Server 2014 is hosted on one Hyper-V failover cluster in your organization's data center. You want to move the guest cluster from its current host Hyper-V failover cluster to a new Hyper-V failover cluster. By migrating one guest cluster node at a time from the original Hyper-V failover cluster to the new Hyper-V failover cluster, you'll be able to continue to service client requests without interruption, failing over SQL Server 2014 to the guest node on the new Hyper-V failover cluster after the first node completes its migration and before migrating the second node across.

Guest host clusters can use one of the following shared storage options:

- **Shared virtual hard disk.** This is a new technology in Windows Server 2012 R2 that allows guest cluster VMs to share a single virtual hard disk file in .vhdx format. This virtual hard disk file functions as shared storage for the guest cluster.
- **Virtual Fibre Channel.** This requires that each Hyper-V failover cluster node is configured with a Fibre Channel Host Bus Adapter that is compatible with Virtual Fibre Channel.
- **iSCSI.** Windows Server 2012 and Windows Server 2012 R2 include iSCSI initiator software. You can use this to connect to an iSCSI target located on your organization's SAN. As Windows Server 2012 and Windows Server 2012 R2 include a role service for an iSCSI target server, it's even possible, though not recommended, to have the iSCSI target running on a Windows Server 2012 or Windows Server 2012 R2 VM.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn440540.aspx>

Shared virtual hard disk

Shared virtual hard disks are a feature new to Windows Server 2012 R2 that allow Hyper-V guest clusters to use specially prepared virtual hard disks as guest cluster shared storage. Each guest cluster node uses a normal virtual hard disk to host the operating system and applications. A shared hard virtual disk must be configured as follows:

- Must be configured to use .vhdx format
- Must be stored on a scale-out file server share or on a Cluster Shared Volume
- Must be connected to each guest cluster VM through a virtual SCSI controller
- Each guest cluster VM must be running the Windows Server 2012 or Windows Server 2012 R2 operating system.
- Must have the Enable virtual hard disk sharing setting configured on each guest cluster VM, as shown in Figure 6i.

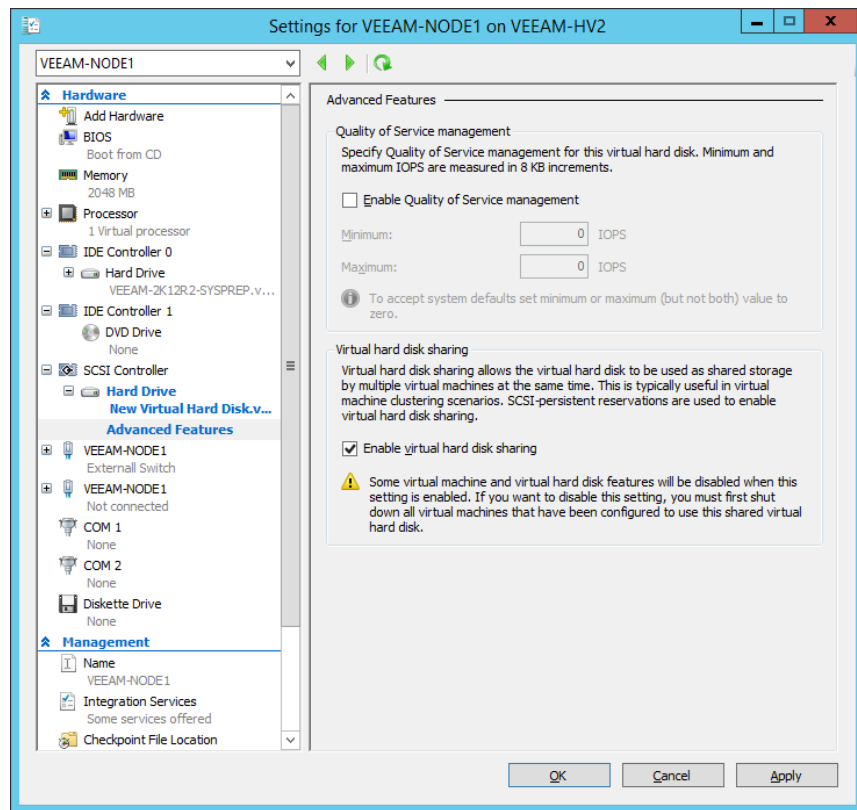


Figure 6i. Enable virtual hard disk sharing

Find Out More:

You can learn more about using Shared Virtual Hard Disks for guest clustering by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/dn265980.aspx>

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. What four technologies can you use to provide shared storage to a Hyper-V failover cluster?
2. What file system format can you use for disk witnesses in Windows Server 2012 R2 failover clusters?]
3. If you deploy a five-node failover cluster using the default settings, what cluster quorum model will be used by default?
4. You are deploying a failover cluster using default settings. What number of nodes, even or odd, would you deploy if you intend to have the cluster use the Node and Disk Majority quorum model? [You would deploy an even number of nodes if you had a disk witness and wanted to deploy a cluster that uses the Node and Disk Majority quorum model.]
5. What feature of Windows Server 2012 R2 failover clusters removes problems related to partitioned clusters?
6. You want to monitor the Cluster Aware Updating process. Which Cluster Aware Updating mode should you use?
7. List three prerequisites of a shared virtual hard disk.

Chapter 7: Virtual Machine Movement

One of the big advantages of virtual machines (VMs) over traditionally deployed servers is their portability. It is fairly straightforward to move a VM between virtualization hosts. You can do this by performing an export and an import of the VM if you don't mind the VM being offline, or by leveraging live migration if you need to move the VM while it is still running and clients are interacting with it.

In this chapter you'll learn about:

- Live migration
- Storage migration
- Exporting, importing, and copying VMs
- VM network health detection
- Drain on shutdown
- P2V migrations
- V2V migrations

Live migration

Live migration is the process of moving an operational VM from one physical virtualization host to another with no interruption to VM clients or users. Live migration is supported between cluster nodes that share storage, between separate Hyper-V virtualization hosts that are not participating in a failover cluster using a SMB 3.0 file share as storage, and even between separate Hyper-V hosts that are not participating in a failover cluster using a process termed "shared nothing live migration."

Live migration has the following prerequisites:

- There must be two or more servers running Hyper-V that use processors from the same manufacturer; for example, all Hyper-V virtualization hosts configured with Intel processors or all Hyper-V virtualization hosts configured with AMD processors.
- Hyper-V virtualization hosts need to be members of the same domain, or domains that have a trust relationship with each other.
- VMs must be configured to use virtual hard disks or virtual Fibre Channel disks (no pass-through disks).

It is possible to perform live migration with VMs configured with pass-through disks under the following conditions:

- VMs are hosted on a Windows Server Hyper-V failover cluster.
- Live migration is going to be within nodes that participate in the same Hyper-V failover cluster.
- VM configuration files are stored on a Cluster Shared Volume.
- The physical disk used as a pass-through disk is configured as a storage disk resource that is controlled by the failover cluster. This disk must be configured as a dependent resource for the highly available VM.

If performing a live migration using shared storage, the following conditions must be met:

- The SMB 3.0 share needs to be configured so that source and destination virtualization host's computer accounts have read and write permissions.
- All VM files (virtual hard disks, configuration and snapshot files) must be located on the SMB 3.0 share. You can use storage migration to move VM files to an SMB 3.0 share while the VM is running prior to performing a live migration using this method.

You must configure the source and destination Hyper-V virtualization hosts to support live migrations by enabling live migrations in Hyper-V settings, as shown in Figure 7a. When you do this, you specify the maximum number of simultaneous live migrations and the networks that you will use for live migration. Microsoft recommends using an isolated network for live migration traffic, though this is not a requirement.

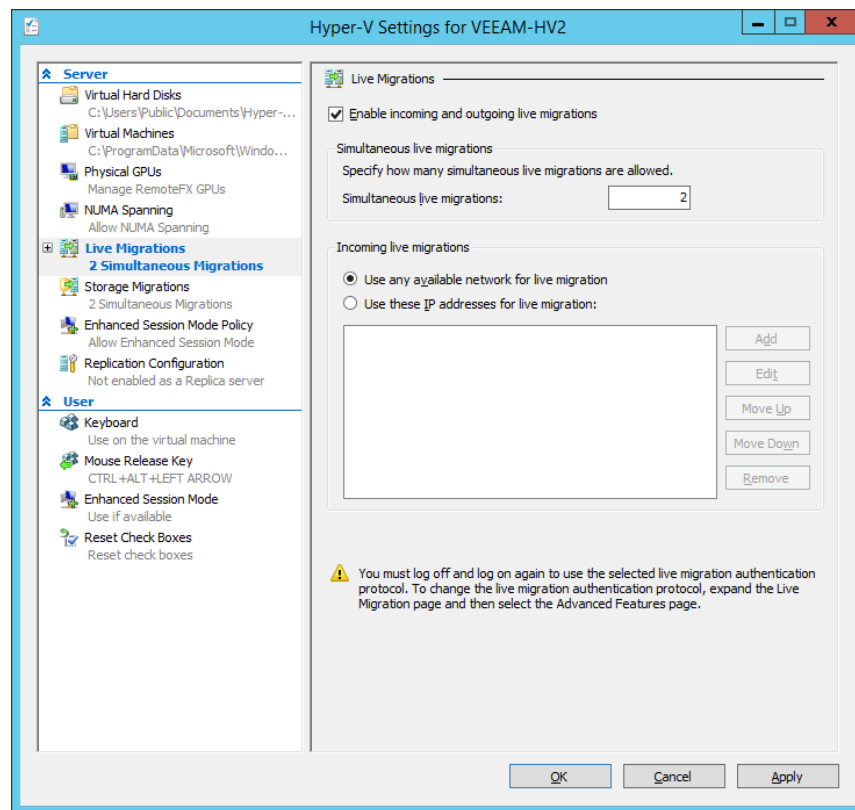


Figure 7a. Enable a live migration

The next step in configuring live migration is choosing which authentication protocol and live migration performance options to use. You select these in the Advanced Features area of the Live Migrations settings. The default authentication protocol is CredSSP (Credential Security Support Provider). CredSSP requires local sign-on to both source and destination Hyper-V virtualization host to perform live migration. Kerberos allows you to trigger live migration remotely. To use Kerberos, you must configure the computer accounts for each Hyper-V virtualization host with constrained delegation for the cifs and Microsoft Virtual System Migration Service services, granting permissions to the virtualization hosts that will participate in the live migration partnership. The performance options allow you to speed up live migration. Compression increases processor utilization. SMB will use SMB Direct if both the network adapters used for the live migration process support Remote Direct Memory Access (RDMA) and RDMA capabilities are enabled. These options are shown in Figure 7b.

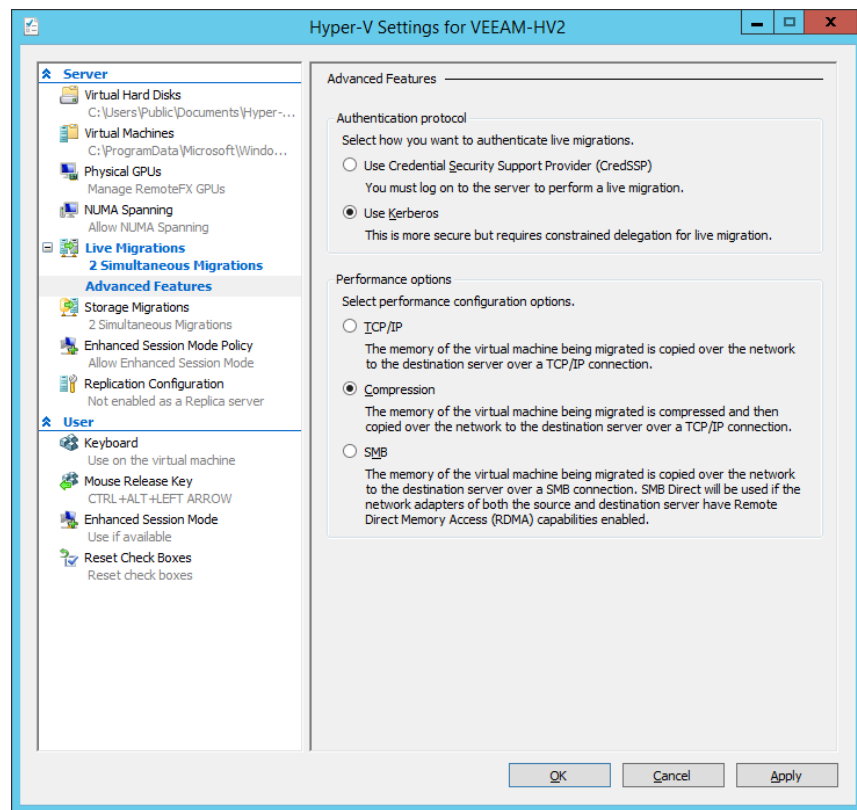


Figure 7b. Live migration authentication and performance

Find Out More:

You can learn more about Live Migration by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831435.aspx>

Storage migration

With storage migration, you can move a VM's virtual hard disk files, checkpoint files, smart paging files, and configuration files from one location to another. You can perform storage migration while the VM is running or while the VM is powered off. You can move data to any location that is accessible to the Hyper-V host. This allows you to move data from one volume to another, from one folder to another, or even to an SMB 3.0 file share on another computer. When performing storage migration, choose the Move the VM's storage option as shown in Figure 7c.

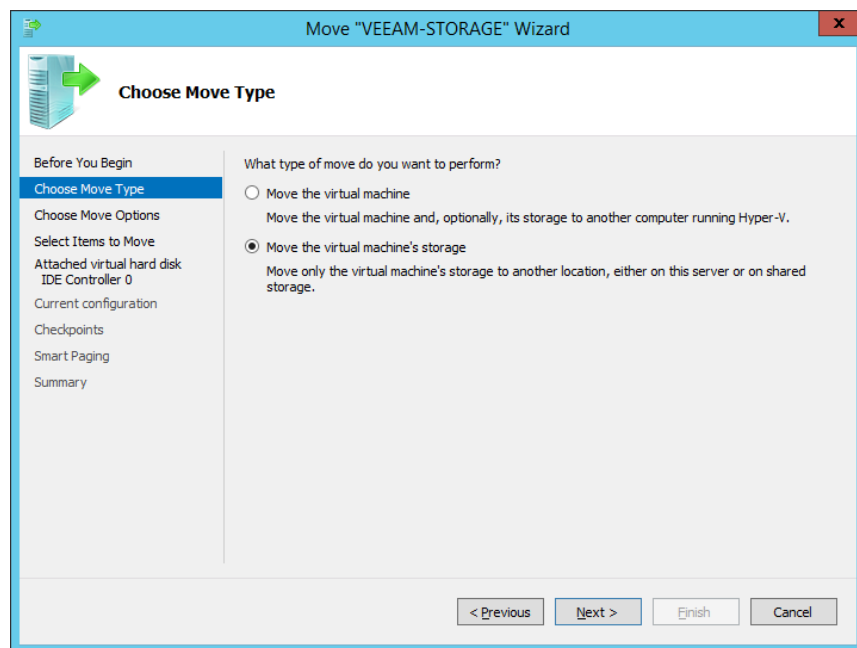


Figure 7c. Move options

For example, you could use storage migration to move VM files from one Cluster Share Volume to another on a Hyper-V failover cluster without interrupting the VM's operation. As figure 7d shows, you have the option of moving all data to a single location, moving VM data to separate locations or moving only the VM's virtual hard disk.

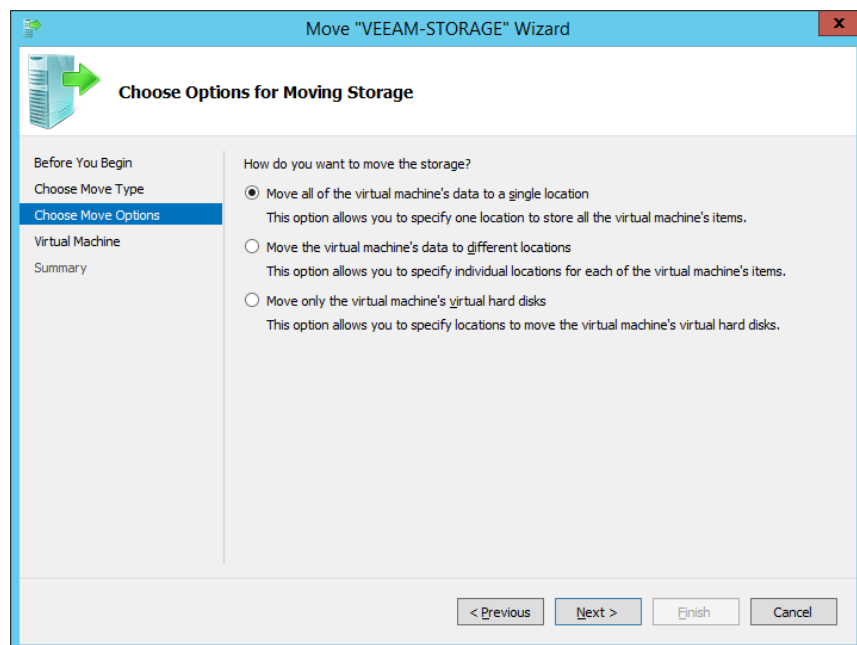


Figure 7d. VM storage move options

To move the VM's data to different locations, select the items you want to move and the destination locations, as shown in Figure 7e.

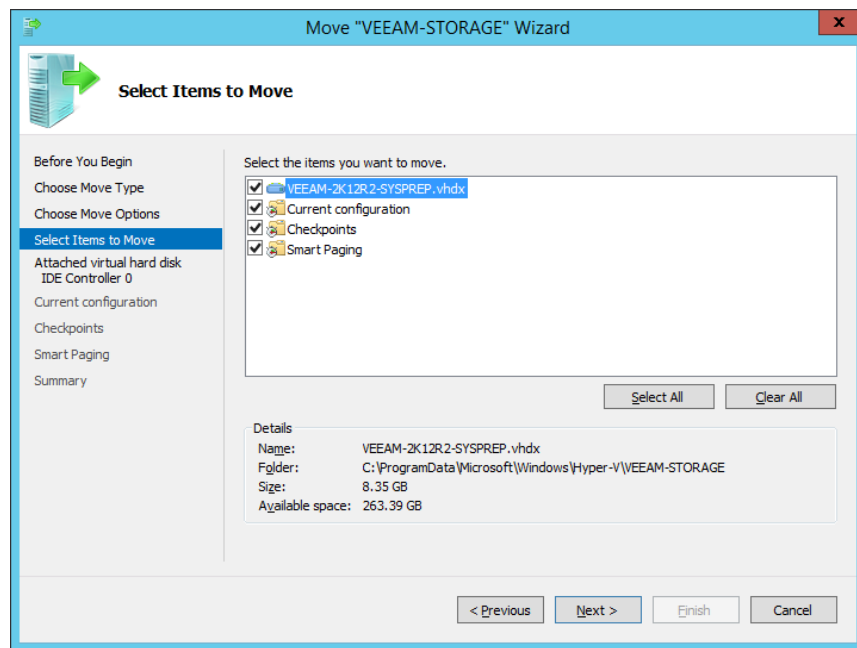


Figure 7e. Select items to move

Find Out More:

You can learn more about storage migration by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh831656.aspx>

Exporting, importing and copying VMs

A VM export creates a duplicate of a VM that you can import on the same, or a different Hyper-V virtualization host. When performing an export, you can choose to export the VM, which includes all its VM checkpoints, or just a single VM checkpoint. Windows Server 2012 R2 supports exporting a running VM. With Hyper-V in Windows Server 2012, Windows Server 2008 R2 and Windows Server 2008, it is necessary to shut down the VM before performing an export.

Exporting a VM with all of its checkpoints will create multiple differencing disks. When you import a VM that was exported with checkpoints, these checkpoints will also be imported. If you import a VM that was running at the time of export, the VM is placed in a saved state. You can resume from this saved state, rather than having to restart the VM.

When importing a VM, choose from the options shown in Figure 7f.

Figure 7f. Import Options

- **Register the virtual machine in place (use the existing ID).** Use this option when you want to import the VM while keeping the VM files their current location. As this method uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which you wish to import the VM.
- **Restore the virtual machine (use the existing unique ID).** Use this option when you want to import the VM while moving the files to a new location; for example, if you are importing a VM that was exported to a network share. As this method also uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which you wish to import the VM.
- **Copy the virtual machine (create a new unique ID).** Use this method if you want to create a separate clone of the exported VM. The exported files will be copied to a new location, leaving the original exported files unaltered. A new VM ID is created, meaning that the cloned VM can run concurrently on the same virtualization host as the original progenitor VM. When importing a cloned VM onto the same virtualization host as the original progenitor VM, ensure that you rename the newly imported VM, otherwise you may confuse them.

Find Out More:

You can learn more about this topic by consulting the following TechNet article:
http://technet.microsoft.com/en-us/library/dn282278.aspx#bkmk_export

VM Network Health Detection

VM Network Health Detection is a feature for VMs that are deployed on Hyper-V host clusters. With VM Network Health Detection, you configure a VM's network adapter settings and mark certain networks as being protected. You do this on the Advanced Features section of the network adapter as shown in Figure 7g.

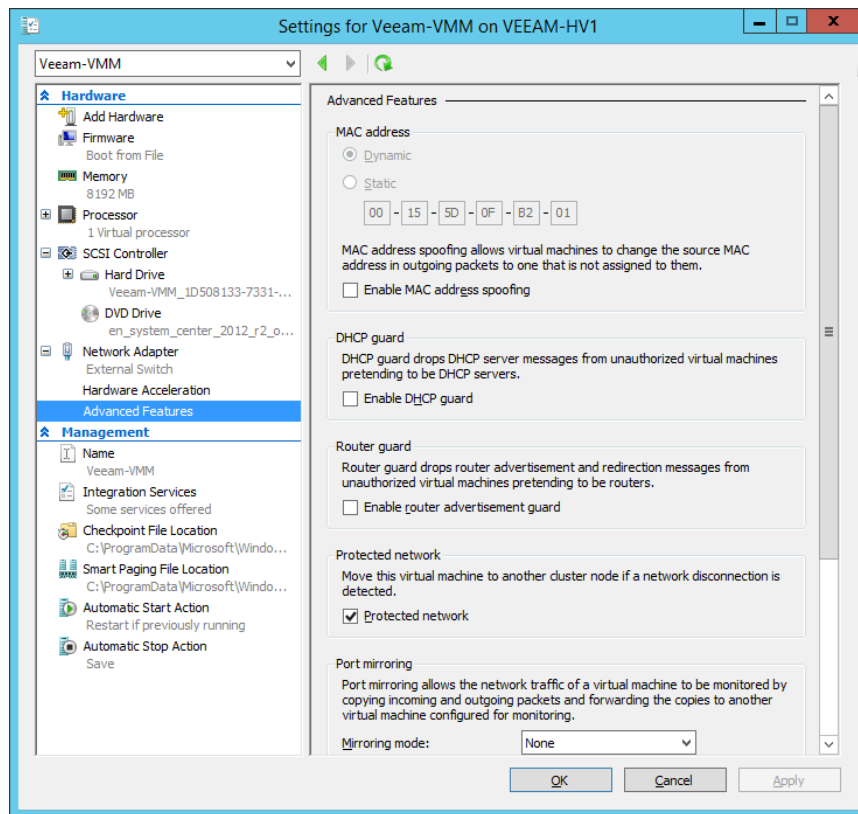


Figure 7g. Protected network option

In the event that a VM is running on a cluster node where the network marked as protected becomes unavailable, the cluster will automatically live migrate the VM to a node where the protected network is available. For example, you have a four-node Hyper-V failover cluster. Each node has multiple network adapters and a virtual switch named Alpha maps as an external virtual switch to a physical network adapter on each node. A VM, configured as highly available and hosted on the first cluster node is connected to virtual switch Alpha. The network adapter on this VM is configured with the protected network option. After the VM has been switched on and has been running for some time, a fault occurs causing the physical network adapter mapped to virtual switch Alpha on the first cluster node to fail. When this happens, the VM will automatically be live migrated to another cluster node where virtual switch Alpha is working.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK_VMHealth

VM drain on shutdown

VM drain on shutdown is a feature new to Windows Server 2012 R2 that will automatically live migrate all running VMs off a node if you shut down that node without putting it into maintenance mode. If you are following best practice, you'll be putting nodes into maintenance mode and live migrating running workloads away from nodes that you will restart or intend to shut down anyway. The main benefit of VM drain on shutdown is that, in the event that you are having a bad day and forget to put a cluster node into maintenance mode before shutting it down or restarting it, any running VMs will be live migrated away without requiring your direct intervention.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK_VMDrain

P2V migrations

P2V migrations allow you to capture and virtualize a physical computer. P2V migrations are not supported in System Center 2012 R2 Virtual Machine Manager (VMM), but are supported in the earlier version, System Center 2012 VMM. One recommended course of action is to use the earlier version of VMM to perform P2V migrations, export the converted VM from the virtualization hosts managed by the earlier version of VMM and then import those converted exported VMs into the virtualization infrastructure managed by System Center 2012 R2 VMM.

Find Out More:

You can learn more about this topic by consulting the following TechNet blog post: <http://blogs.technet.com/b/scvmm/archive/2013/10/03/how-to-perform-a-p2v-in-a-scvmm-2012-r2-environment.aspx>

V2V migrations

A V2V migration occurs when you migrate a VM from one virtualization host, for example VMware's ESXi platform, so that the VM is hosted on Hyper-V. System Center 2012 R2 VMM supports V2V migration from VMs running on the following versions of VMware:

- VMware ESX 4.1
- VMware ESXi 4.1
- VMware ESXi 5.1

Prior to performing a V2V conversion of a VM hosted on a supported version of VMware, you must uninstall VMware tools from the host OS of the VM to be converted, disable anti-malware software running within the host OS of the VM to be converted, and power off the VM to be converted. You also can't use VMM to convert a VMware VM that uses virtual hard disks connected to an IDE bus.

Find Out More:

You can learn more about this topic by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/gg610672.aspx>

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. You want to be able to use SMB direct with VM live migration. What technology must be supported and enabled to accomplish this goal?
2. For which services must you configure constrained delegation on the source and destination virtualization host computer accounts if you are using Kerberos to authenticate live migration traffic?
3. You recently exported a VM named SYD-VM1 from host SYD-HV1. You now want to import SYD-VM1 onto host SYD-HV1, renaming it SYD-VM1-DEV. You wish to keep SYD-VM1 on SYD-HV1. Which VM import option should you select?
4. You have a VM hosted on a Windows Server 2012 R2 Hyper-V failover cluster. The VM is configured to use an external switch. You want the VM to live migrate to another cluster node if the external switch loses connectivity on its current node, but connectivity for the virtual switch is available on another node. Which feature should you enable?

5. You want to perform storage migration to a file share hosted on another computer. Assuming the share permissions are configured correctly, what technology must the file share support to be able to host the Hyper-V configuration, checkpoint, virtual hard disk and smart paging files?
6. Which Windows Server 2012 R2 feature ensures that running VMs are live migrated off a cluster node if the node is shut down or restarted without being put into maintenance mode?
7. Which version of VMM does not support P2V migrations?
8. What steps must you take before performing a V2V migration of a VM that uses virtual hard disks attached to virtual SCSI adapters hosted on VMware ESXi 5.1 using VMM to a Hyper-V host running on Windows Server 2012 R2?

Chapter 8: Monitoring and Hyper-V Replica

Monitoring is a critical component of any enterprise-scale deployment. Unless you perform effective monitoring, you won't be aware of how well your virtualization infrastructure is performing or if any faults that are worthy of your notice have occurred.

Hyper-V replica is a technology new to Windows Server 2012 that allows you to configure a replica virtual machine (VM) on a second Hyper-V host that functions as a lagged copy of a VM running on the primary host.

In this chapter you'll learn about:

- Audit Collection Services
- System Center Global Service Monitor
- Monitoring fabrics
- Operations Manager reporting
- Operations Manager management packs
- Monitoring Active Directory
- Domain controller cloning
- Hyper-V replica
- Hyper-V Recovery Manager

Audit Collection Services

Audit Collection Services is an additional Operations Manager component that allows you to centralize the collection of security log data from computers within your organization. For example, you configure file and folder auditing on all of the file servers in your organization as a way of tracking file and folder access. Rather than trawling through the security event logs on each file server, Audit Collection Services gives you a central location that you can use to search and analyze this data.

Audit Collection Services can be deployed separate from Operations Manager. This allows you to separate auditing data from other data recorded by Operations Manager. The Microsoft Audit Collection Services node in the Monitoring Workspace of the Operations Manager console is shown in Figure 8a.

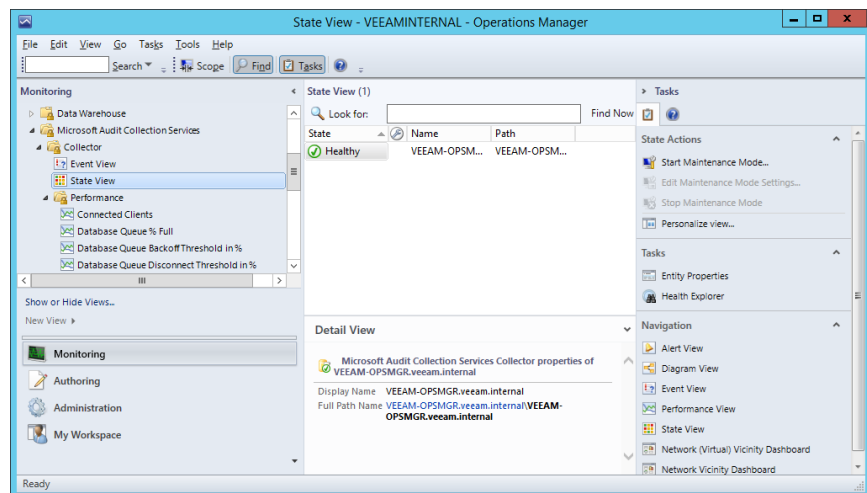


Figure 8a. Audit Collection Services state view

You can use ACS to collect security auditing information from computers running Windows Server operating systems, Oracle Solaris, IBM AIX, and supported UNIX flavors and Linux distributions.

An ACS deployment includes the following components:

- **ACS collector.** Deployed on a server, this component accepts incoming event log data, transferring it to the ACS database.
- **ACS database.** A SQL Server database that stores event log data collected by ACS.
- **ACS forwarder.** The ACS forwarder is a service that runs on computers from which you want to collect event log data. The forwarder extracts relevant data from the event logs and transmits it to the ACS collector.
- **ACS reporting server.** The ACS reporting server uses SQL Server Reporting Services (SSRS) to generate reports using auditing data. The ACS reporting server can use the SSRS instance that supports the Operations Manager deployment, or it can use a separate instance of SSRS.

Find Out More:

You can learn more about Audit Collection Services by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh212908.aspx>

System Center Global Service Monitor

System Center Global Service Monitor is a cloud service available from Microsoft that you can use to monitor whether web-based applications are accessible by testing them from a variety of locations around the world. For example, System Center Global Service Monitor can provide you with web application availability information by simulating customer traffic from Australia, Switzerland, Russia, Brazil and the United States. System Center Global Service Monitor interacts with the application, providing more detailed availability information than a simple test of whether the service hosting the application responds to HTTP requests.

System Center Global Service Monitor provides a general Web Application Availability monitor that checks the responsiveness of specific URLs. System Center Global Service Monitor also provides Visual Studio Web Tests, which allow you to run complex, multiple step, authenticated procedures. Tests can be run every five minutes.

You can use System Center Global Service Monitor with the System Center 2012 R2 Operations Manager console. When you do this, System Center Global Service Monitor appears as a node in the Administration

Find Out More:

You can learn more about System Center Global Service Monitor by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj860368.aspx>

Fabric monitoring

By integrating System Center 2012 R2 Virtual Machine Manager (VMM) with Systems Center 2012 R2 Operations Manager and utilizing the management pack for VMM, you can monitor the availability not only of VMM, but also the availability, health and performance of all VMs and virtualization hosts managed by VMM. You can use diagram views to view all components in a fabric, allowing you to quickly determine which components require attention. Figure 8b shows a diagram view for the all of the components managed by a VMM management server.

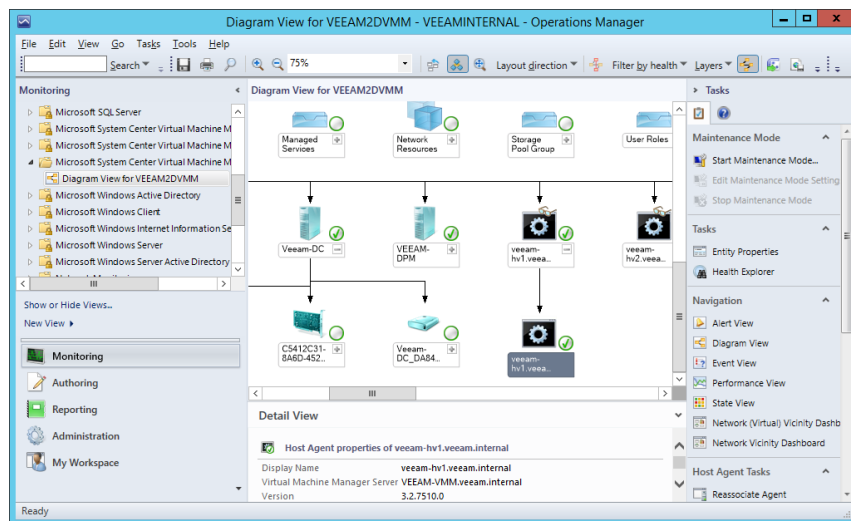


Figure 8b. Diagram view

The Fabric Health Dashboard provides information about the health of each VMM private cloud as well as the fabric that supports those clouds. For each cloud, the Fabric Health Dashboard provides the following information:

- **Compute.** Allows you to view hosts, operating systems running on hosts, workloads, workload configuration and management infrastructure.
- **Storage.** Displays file shares, LUNs and Storage Pools
- **Hardware.** Displays physical and virtual network devices, networks and fabric hardware.

Find Out More:

You can learn more about fabric monitoring by consulting the following TechNet <http://technet.microsoft.com/en-us/library/dn458592.aspx>

Operations Manager reporting

Operations Manager can use SQL Server Reporting Services (SSRS) to generate complex reports based on data stored in the Operations Manager data warehouse. This allows you to view recently collected data as well as configure reports to view historical trends, such as growth in virtualization resource utilization. You can increase the number of available reports, either by importing management packs or by authoring your own reports. Figure 8c shows the reports available from the Active Directory Server Common Library.

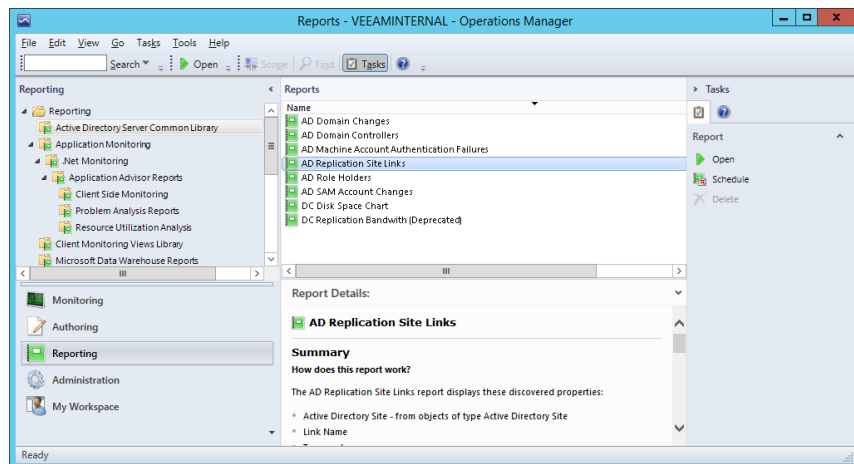


Figure 8c. Reporting workspace

Reporting is an optional component and requires the deployment of SSRS. Other applications, including other System Center 2012 R2 components, are unable to use the SSRS instance that is utilized by Operations Manager reporting.

When running a report, select the time period you want the report to encompass as well as the objects relevant to the report. Figure 8d shows the setup for a virtualization host utilization report for servers veeam-hv1 and veeam-hv2.

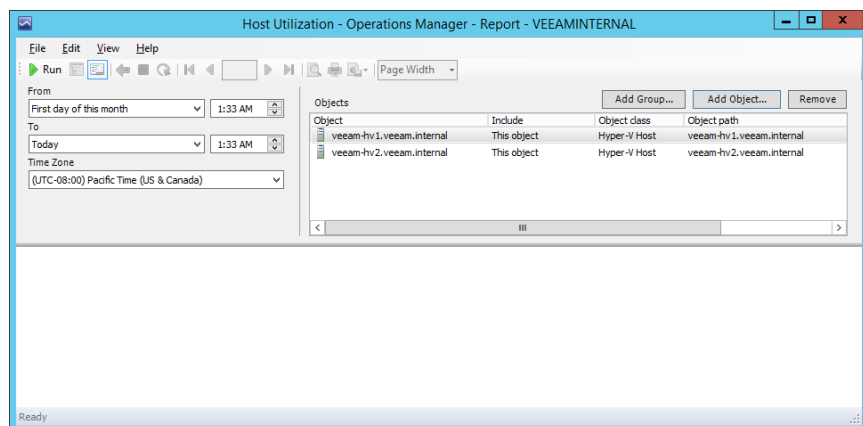


Figure 8d. Virtualization host report configuration

You can schedule reports to run on a periodic basis. Reports can be run once or on an hourly, daily, weekly or monthly basis. Figure 8e shows configuring the schedule for a report subscription.

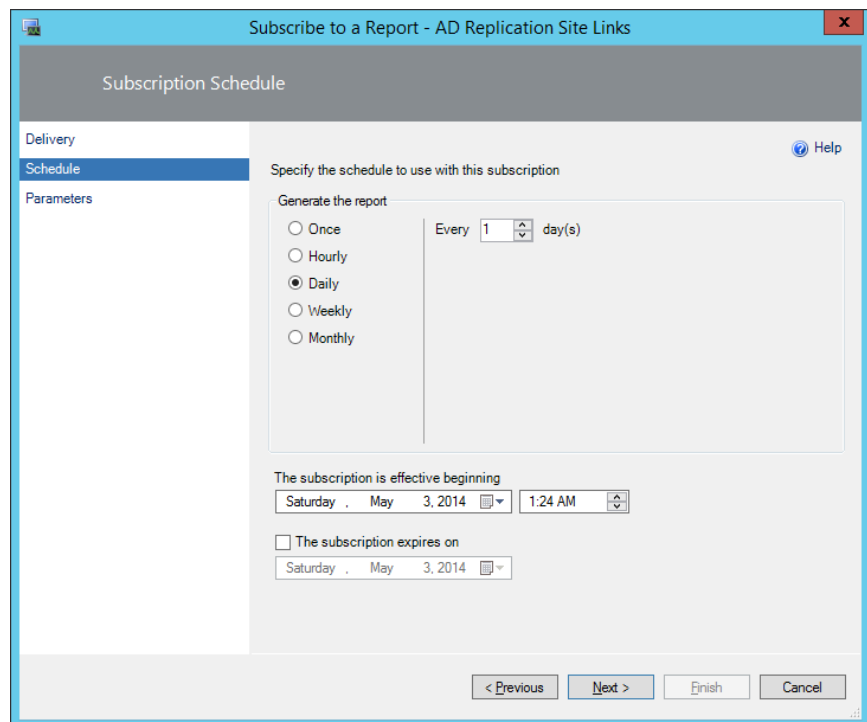


Figure 8e. Report schedule

Find Out More:

You can learn more about Operations Manager reporting by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh212786.aspx>

Operations Manager management packs

Management packs allow you to extend the functionality of Operations Manager by improving the product's ability to interact with specific applications or services. Management packs are usually created for specific products and reflect advice about the performance and configuration of that product, either by the organization that created the product, or by groups that have expert knowledge about the product's performance characteristics and functionality. Figure 8f shows the Management Packs node of the Operations Manager console.

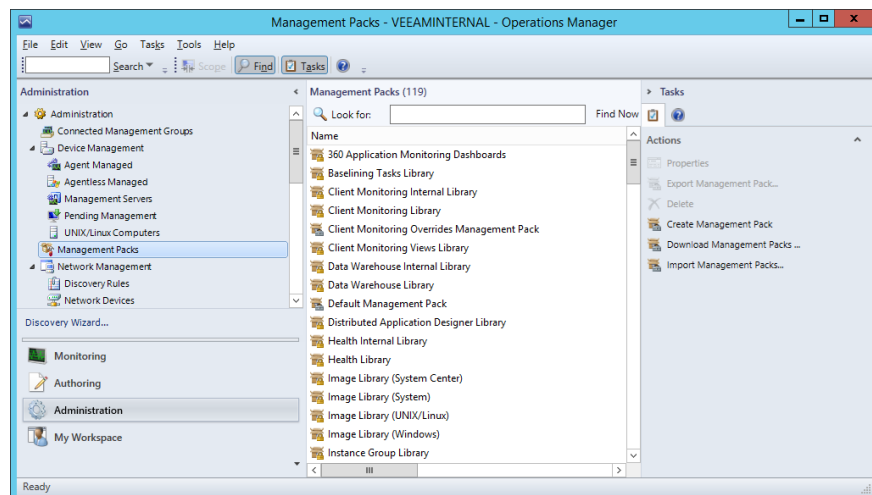


Figure 8f. Management packs

A management pack can consist of some or all of the following components:

- **Monitors.** These direct the Operations Manager agent to track the state of managed components. For example, a monitor might track the length of a message queue on an Exchange server. This monitor may have three states: green if there are less than 20 messages in the queue, yellow if there are between 20 and 50 messages in the queue and red if there are more than 50 messages in the queue.
- **Rules.** Rules determine how an agent collects performance and event data as well as what events trigger Operations Manager alerts. Rules can also run scripts.
- **Tasks.** Tasks are activities that can be executed by the Operations Manager agent or the console; for example, a task that restarts a service or applications or a task that runs a process such as deduplication.
- **Knowledge.** Articles and written advice related to the product, service or application covered by the management pack.
- **Views.** Provides custom monitoring and management interfaces.
- **Reports.** Includes specialized methods of providing information about the component.
- **Object discoveries.** Allows the identification of objects and components that can be monitored.
- **Run As profiles.** Inclusion of profiles that allow rules, tasks, monitors and discoveries to be run using an alternate set of credentials

Management packs are either in a sealed state or an unsealed state. A sealed management pack is a digitally signed binary file that can't be modified. The majority of management packs that you will obtain from third-party vendors or Microsoft will be sealed as a way of validating their integrity. Sealed management packs use the .mp(b) file extension. You can customize a sealed management pack by configuring overrides or creating additional rules, monitors and tasks. You save these customizations in an additional management pack file. Unsealed management packs are in .xml format and allow you to modify settings.

Library management packs are a special type of management pack that stores classes that other management packs may require. A dependency exists when one management pack requires components stored in another management pack. When installing management packs, you must ensure that the management pack dependencies have been met. You can view a management pack's dependencies by viewing the management pack's properties as shown in Figure 8g.

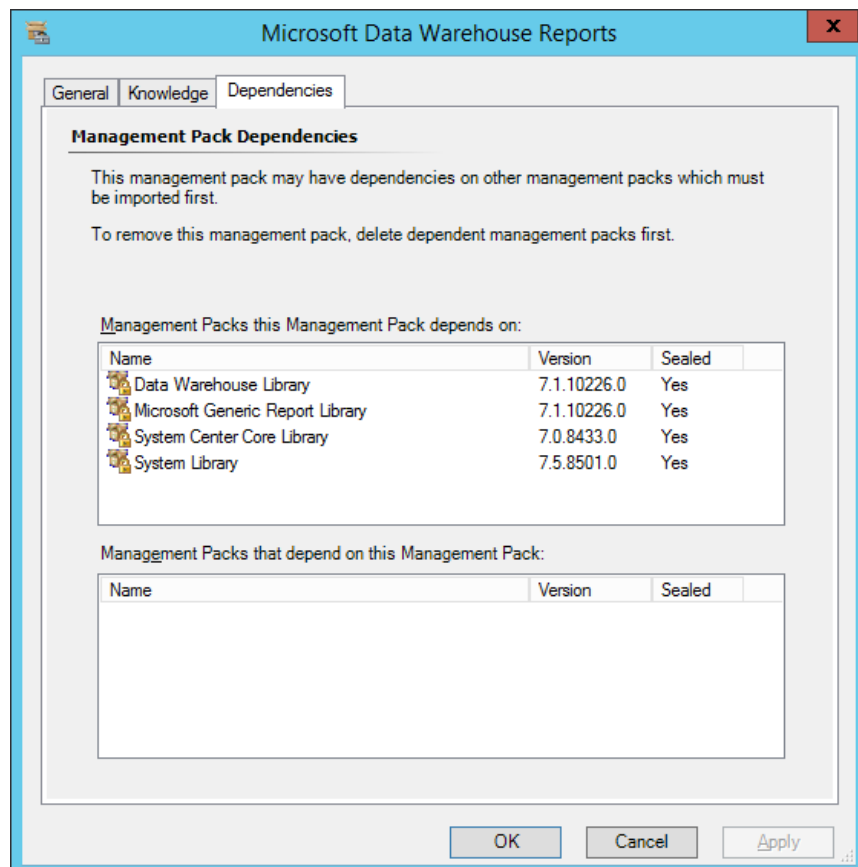


Figure 8g. Caption

Microsoft publishes management packs to an online management pack catalog. You can connect to this management pack catalog from the Operations Manager console and use it to download these management packs as shown in Figure 8h. You can also import management packs from other locations, such as management packs that you downloaded from a specific vendor's website.

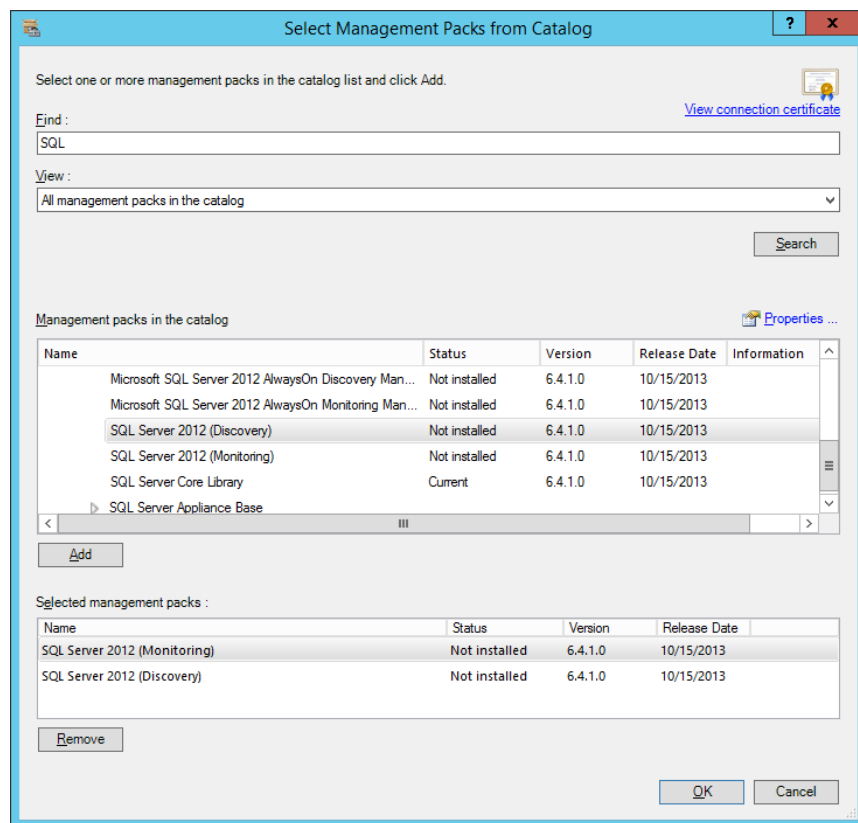


Figure 8h. Caption

Find Out More:

You can learn more about Operations Manager management packs by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/hh212794.aspx>

Monitoring Active Directory

The Active Directory Domain Services Management Pack is a management pack designed for System Center 2012 R2 Operations Manager that you can download free from Microsoft's website. When you have deployed Operations Manager agents to your organization's domain controllers and installed this management pack, you can monitor the performance and availability of domain controllers including:

- Replication
- Lightweight Directory Access Protocol (LDAP)
- Domain Controller Locator
- Trusts
- Net Logon service
- File Replication Service (FRS)
- Intersite Messaging service
- Windows Time service
- Active Directory Web Services (ADWS)
- Active Directory Management Gateway Service
- Key Distribution Center (KDC)
- Service availability
- Key performance data
- Generate reports about service availability and health

The AD Forests and Domains Topology View from this management pack is shown in Figure 8i.

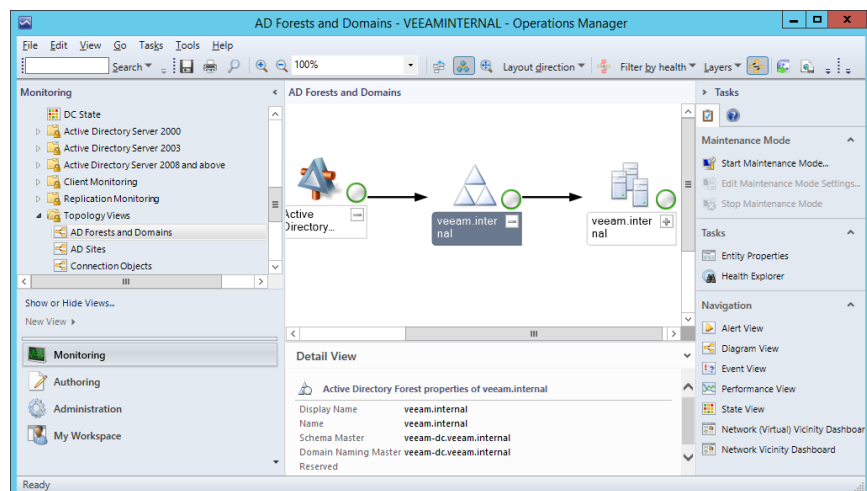


Figure 8i. Caption

Find Out More:

You can learn more about the Active Directory Domain Services Management Pack by consulting the following TechNet article: <http://www.microsoft.com/en-au/download/details.aspx?id=21357>

Domain controller cloning

Windows Server 2012 and Windows Server 2012 R2 support creating copies of domain controllers that are running as VMs as long as certain conditions are met. Cloned domain controllers have the following prerequisites:

- The virtualization host supports VM-GenerationID, a 128-bit random integer that identifies each VM checkpoint. The version of Hyper-V available with Windows Server 2012 and later, as well as some third-party hypervisors, support VM-GenerationID.
- The domain controller must be running Windows Server 2012 or later as its operating system.
- The server that hosts PDC emulator Flexible Single Master Operations Role (FSMO) must be able to be contacted. This server must be running Windows Server 2012 or later as its operating system.
- The computer account of the domain controller that will serve as the template for cloning must be added to the Cloneable Domain Controllers security group as shown in Figure 8j.

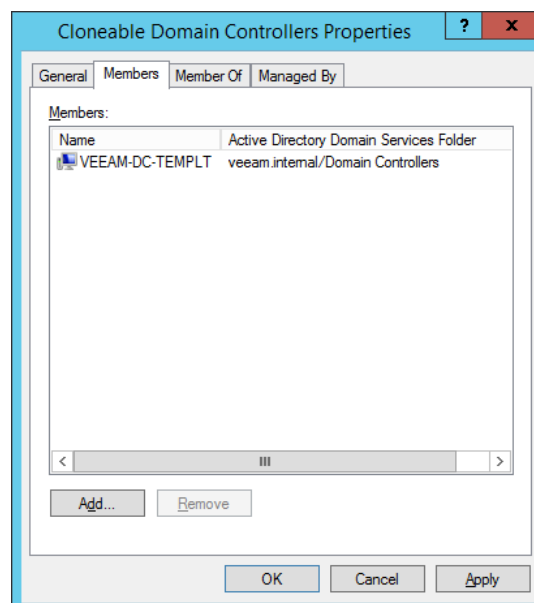


Figure 8j. Caption

Once you have met these conditions, you'll need to create an XML configuration file named `DCCloneConfig.xml` using the **New-ADDCCloneConfig** Windows PowerShell cmdlet. Once created, you'll need to edit this file and specify settings such as computer name, network settings and Active Directory site information. You should also check the template DC using the **Get-ADDCCloneExcludedApplicationsList** cmdlet to determine if any services that will cause problems with the cloning, such as the DHCP server service, are present on the template DC.

Find Out More:

You can learn more about domain controller cloning by consulting the following TechNet article: <http://blogs.technet.com/b/askpfeplat/archive/2012/10/01/virtual-domain-controller-cloning-in-windows-server-2012.aspx>

Hyper-V replica

Hyper-V replica is a feature new to Windows Server 2012 that allows you to configure a replica VM running on a separate Hyper-V host. The replica is a lagged but consistent copy of the original. This means that while the replica won't be a precise clone of the original, it will be a consistent copy of the original as it existed for a period up to several minutes in the past. For example, you have two separate Hyper-V virtualization hosts, HV1 and HV2. HV1 hosts a VM named VM1. You can configure Hyper-V replica so that a replica of VM1 is hosted on HV2. Depending on when replication occurs, the replica of VM1 on HV2 will represent the state of VM1 on HV1 as it existed within the last few minutes. The replicas will be in a consistent state. This means that you will be able to start the replica without worrying that some components are out of sync with others.

Hyper-V replica is an option that allows site-level fault tolerance for VMs. To achieve this goal, you configure replicas of VMs hosted in a production site at a DR (Disaster Recovery) site.

A feature new to Windows Server 2012 R2 is Hyper-V extended replica. Hyper-V extended replica allows you to configure a second replica using the VM on the replica server. For example in the scenario of VM1 hosted on HV1 and replicated by Hyper-V replica to HV2, you could configure virtualization HV3 to host a replica of the replica of VM1 which is located on HV2.

Find Out More:

You can learn more about Hyper-V replica by consulting the following TechNet article: <http://blogs.technet.com/b/yungchou/archive/2013/01/10/hyper-v-replica-explained.aspx>

Configuring Hyper-V Replica Servers

To configure Hyper-V replication, you need to configure the Replication Configuration settings as shown in Figure 8k:

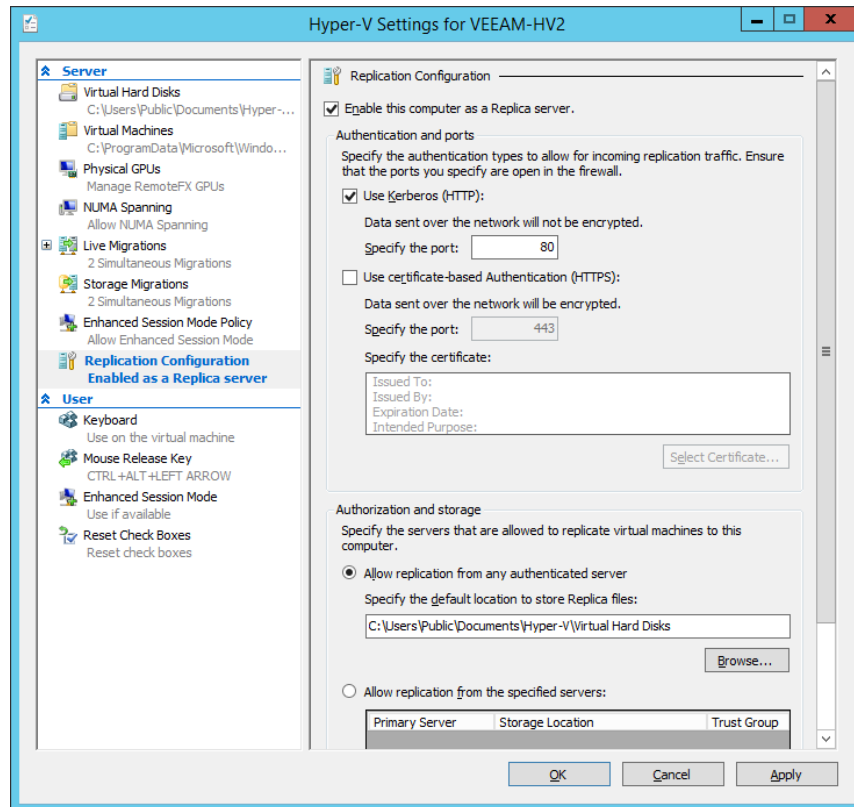


Figure 8k. Caption

When configuring the settings on this dialog box, you first need to select the checkbox for Enable this computer as a Replica server. Next select the authentication method you are going to use. If the computers are parts of the same Active Directory environment, you can use Kerberos. When you use Kerberos, Hyper-V replication data isn't encrypted when transmitted across the network. If you are concerned about encrypting network data, you could configure IPsec. Another option if you are concerned about encrypting replication traffic, which is useful if you are transmitting data across the public internet without using an encrypted VPN tunnel, is to use certificate-based authentication. When using certificate-based authentication, you'll need to import and select a public certificate issued to the partner server using this dialog.

The final step when configuring Hyper-V replica is to select the servers from which the Hyper-V virtualization host will accept incoming replicated VM data. One option is to choose for the Hyper-V virtualization host to accept replicated VMs from any authenticated Hyper-V virtualization host, using a single default location to store replica data. The other option is to configure VM replica storage on a per server basis. For example, if you wanted to store VM replicas from one server on one volume and VM replicas from another server on a different volume, you'd use the second option.

Once replication is configured on the source and destination servers, you'll also need to enable the pre-defined firewall rules to allow the incoming replication traffic. There are two rules: one for replication using Kerberos (HTTP) on port 80 and the other for using certificate-based authentication on port 443.

Configure VM replicas

Once you have configured the source and destination replica servers, you need to configure replication on a per VM basis. You do this by running the Enable Replication wizard, which you can trigger by clicking Enable Replication when the VM is selected in Hyper-V Manager. To configure VM replicas, you must perform the following steps:

- **Select Replica Server.** Select the replica server name. If you are replicating to a Hyper-V failover cluster, you'll need to specify the name of the Hyper-V replica broker. You'll learn more about Hyper-V replica broker later in this chapter.
- **Choose Connection Parameters.** Specify the connection parameters. The options will depend on the configuration of the replica servers. On this page, depending on the existing configuration, you can choose the authentication type and whether replication data will be compressed when transmitted over the network. This page of the wizard is shown in Figure 81.

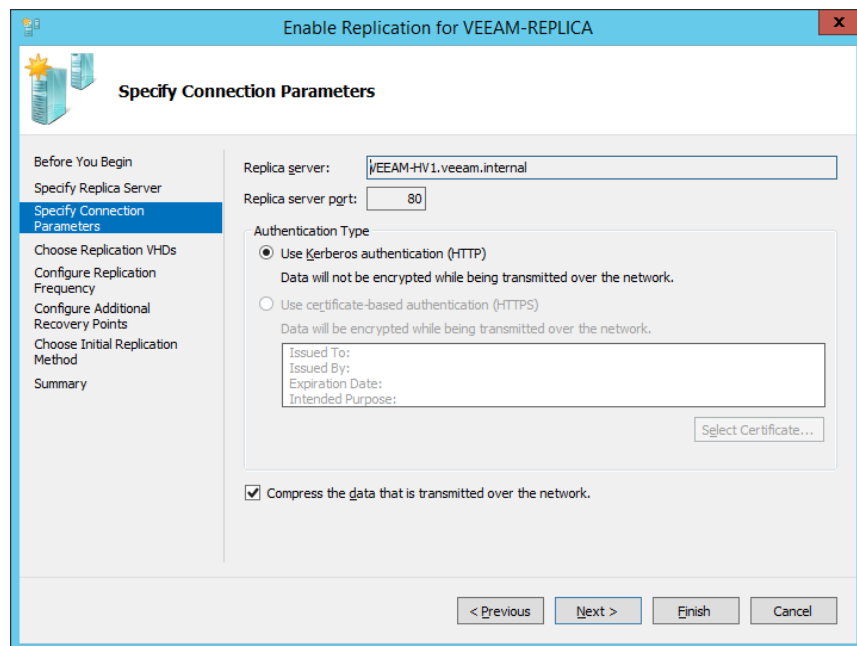


Figure 81. Caption

- **Select Replication VHDs.** When configuring replication, you have the option of not replicating some of a VM's virtual hard disks. In most scenarios, you should replicate all of a VM's hard disk drives. One reason not to replicate a VM's virtual hard disk would be if the virtual hard disk only stores frequently changing temporary data that wouldn't be required when recovering the VM.
- **Replication Frequency.** Use this to specify the frequency with which changes are sent to the replica server. With Windows Server 2012 R2, you can choose between intervals of 30 seconds, 5 minutes and 15 minutes.
- **Additional Recovery Points.** You can choose to create additional hourly recovery points. Doing this gives you the option of starting the replica from a previous point in time rather than the most recent. The advantage is that this allows you to roll back to a previous version of the VM in the event that data corruption occurs and has replicated to the most recent recovery point. The replica server can store a maximum of 24 recovery points. Figure 8m shows the Configuring Additional Recovery Points page.

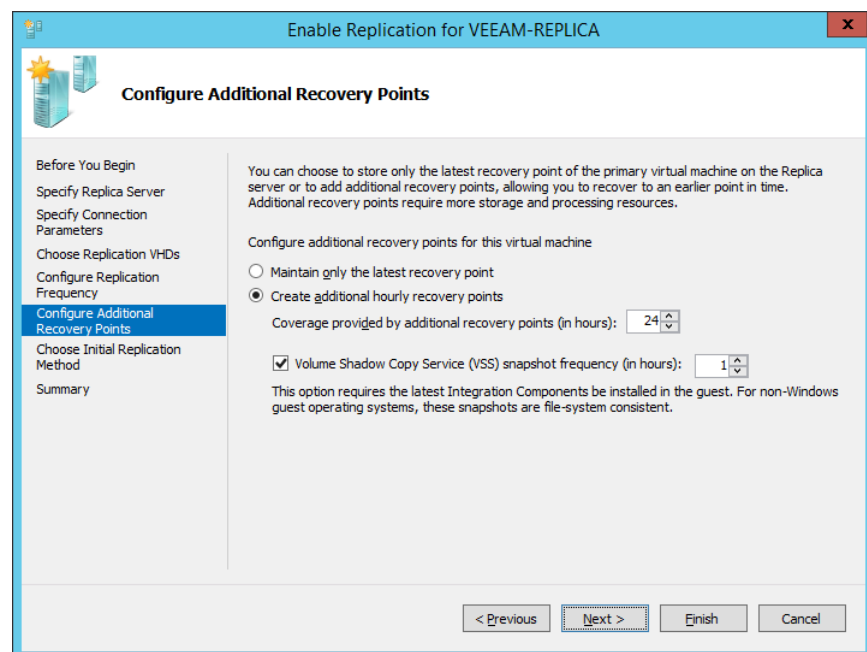


Figure 8m. Caption

- **Initial Replication.** The last step in configuring Hyper-V replica is choosing how to seed the initial replica. Replication works by sending changed blocks of data, so the initial replica, which sends the entire VM, will be the largest transfer. You can perform an offline transfer with external media, use an existing VM on the replica server as the initial copy (the VM for which you are configuring a replica must have been exported and then imported on the replica server), or transfer all VM data across the network. You can perform replication immediately or at a specific time in the future, such as 2 a.m. when network utilization is lower. Figure 8n shows this dialog box.

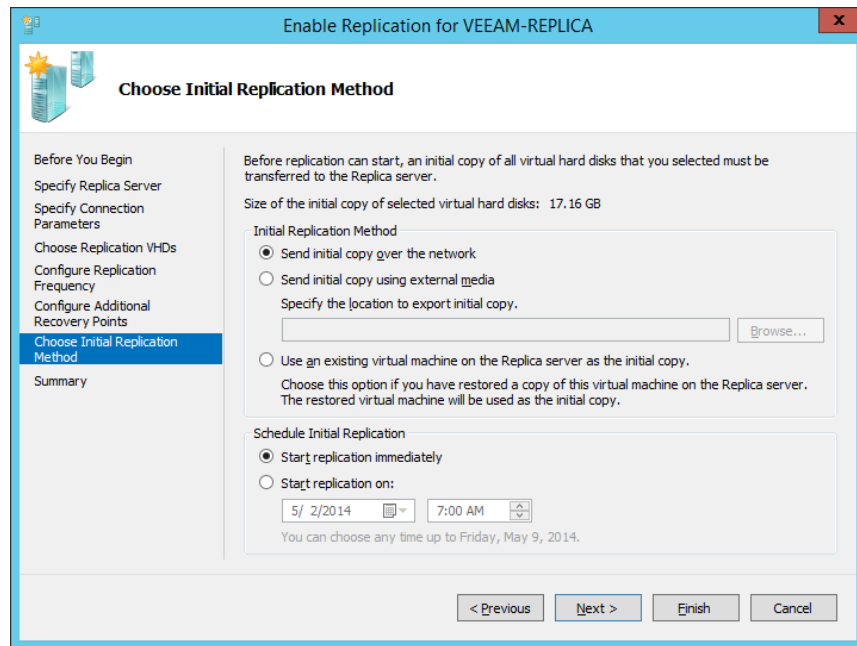


Figure 8n. Caption

Replica failover

You perform planned replica failover when you want to run the VM on the replica server rather than on the primary host. Planned failover involves shutting down the VM, which ensures that the replica will be up to date. Contrast this with Hyper-V live migration, which you perform while the VM is running. When performing planned failover, you can configure the VM on the replica server to automatically start once the process completes and to configure reverse replication, so that the current replica server becomes the new primary and the current primary becomes the new replica server as shown in Figure 8o.

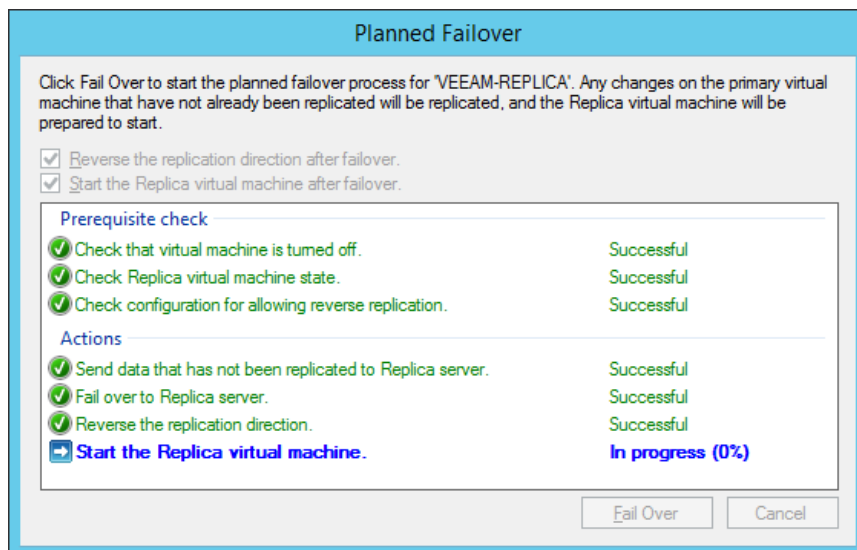


Figure 8o. Caption

In the event that the primary server becomes unavailable, you can trigger an unplanned failover. You would then perform the unplanned failover on the replica server (as the primary is not available). As Figure 8p shows, when performing an unplanned failover, you can select any of the up to 24 previously stored recovery points.

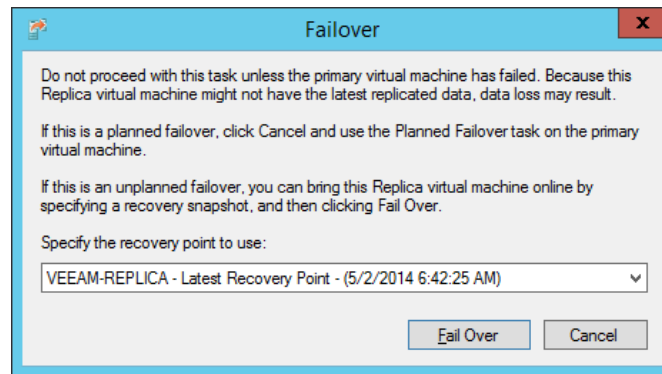


Figure 8p. Caption

Find Out More:

You can learn more about Planned Failover by consulting the following TechNet article: <http://technet.microsoft.com/en-us/library/jj134194.aspx>

Hyper-V replica broker

You need to configure and deploy Hyper-V replica broker if your Hyper-V replica configuration includes a Hyper-V failover cluster as a source or destination. You don't need to configure and deploy Hyper-V replica broker if both the source and destination servers are not participating in a Hyper-V failover cluster. You install the Hyper-V Replica Broker role using Failover Cluster manager after you've enabled the Hyper-V role on cluster nodes.

Find Out More:

You can learn more about Hyper-V replica broker by consulting the following TechNet article: <http://blogs.technet.com/b/virtualization/archive/2012/03/27/why-is-the-quot-hyper-v-replica-broker-quot-required.aspx>

Hyper-V Recovery Manager

Hyper-V Recovery Manager is a service that is located in Microsoft's Windows Azure cloud that manages Hyper-V replication for Virtual Machine Manager VMs. With Hyper-V replica, you have to perform failover manually. Hyper-V Recovery Manager allows you to automate that process. The three key components of Hyper-V Recovery Manager are as follows:

- Automates protection by managing on-going VM replication to a secondary site.
- Performs continuous health monitoring to determine service availability at each monitored site.
- Performs orchestrated recovery in the event of a service outage at your primary site. You configure a recovery plan within Hyper-V Recovery Manager that specifies how recovery at the secondary site occurs.

Find Out More:

You can learn more about Hyper-V Recovery Manager by consulting the following TechNet article: <http://blogs.technet.com/b/windowsserver/archive/2014/01/16/windows-azure-hyper-v-recovery-manager-is-now-generally-available.aspx>

Review

The following set of questions test your understanding of the content of this chapter. Answers are located in the appendix.

1. What file extension do sealed management packs use?
2. Besides the database engine, which SQL Server component must be installed before you can configure and deploy reporting in Operations Manager?
3. Which Operations Manager feature allows you to collect and analyze information about which users accessed files and folders?
4. Which dashboard provides you with information about the health of each VMM private cloud as well as the fabric that supports those clouds?
5. Which security group must a virtualized domain controller be a member of if you wish to use it as a template to create clones?
6. Which authentication method should you choose for Hyper-V replica if you want data sent over the network in an unencrypted format on port 80?
7. What is the maximum number of recovery points you can configure for a Hyper-V replica?
8. In what state must a VM be on the primary Hyper-V host before you can perform a planned failover to the replica host?
9. Which component of the System Center Suite offers monitoring of external web applications from multiple locations around the world?

Appendix

Chapter 1

1. *Smart paging is a Hyper-V feature that allows VMs to restart when their required minimum amount of memory is available, but the amount configured as Startup RAM is not.*
2. *You can decrease, but not increase, the Minimum RAM value and increase, but not decrease, the Maximum RAM value while the VM is running.*
3. *You must enable the guest services integration service to use the **Copy-VMFile** Windows PowerShell cmdlet to copy a file from the virtualization host to the VM.*
4. *You would use the **Enable-VMResourceMetering** Windows PowerShell cmdlet to configure a Hyper-V virtualization host to record CPU, disk, memory and network utilization on a per-VM basis.*
5. *Only generation 2 VMs can boot off a virtual hard disk attached to a virtual SCSI controller. Only VMs running Windows 8, Windows 8.1, Windows Server 2012 and Windows Server 2012 R2 as a guest operating system can be configured as generation 2 VMs.*
6. *Ensure that the user's account on the VM is a member of the Remote Desktop Users local security group.*

Chapter 2

1. *You would use the .vhdx virtual hard disk format if you wanted to provision a VM with a 3072-GB boot volume.*
2. *You need to have reduced the size of the volume using the VM's OS. The VM must be hosted on a computer running Windows Server 2012 R2 Hyper-V. The virtual hard disk needs to use .vhdx format. The virtual hard disk must be connected to a virtual SCSI controller.*
3. **Restore-VMSnapshot.** *The cmdlet has yet to be updated to use the new checkpoint naming system.*
4. *You use the **New-VHD** Windows PowerShell cmdlet to create a differencing virtual hard disk.*
5. *The host computer needs to have a Fibre Channel host bus adapter (HBA) that has a driver that supports virtual Fibre Channel.*
6. *You would configure Storage QoS to limit each virtual hard disk's maximum number of IOPS.*

7. *The disk must be in an offline state to be connected as pass-through storage for a Hyper-V VM.*
8. *Windows Server 2012 and Windows Server 2012 R2.*
9. *The shared virtual hard disk needs to be in .vhdx format, connected to a VM's SCSI controller, not configured as the boot volume, and the VMs must be running Windows Server 2012 or Windows Server 2012 R2.*
10. *You will need to create a dynamically expanding disk that uses the .vhdx virtual hard disk format. The disk needs to be dynamically expanding rather than fixed because you only have access to 4 TB at the moment. It needs to be .vhdx rather than .vhd, as .vhd format virtual hard disks are limited to 2040 GB in size.*

Chapter 3

1. *You would configure an internal Hyper-V virtual switch as this allows network communication between the Hyper-V host and the VM without allowing the VM access to the external network.*
2. *Configure the maximum bandwidth management setting on the VM's virtual network adapter.*
3. *You'll need to configure the VM's network adapter with a static MAC address.*
4. *You can configure the virtual network adapters connected to each VM to use a different VLAN ID. This will separate the traffic from each VM into a different broadcast domain.*
5. *You will need to configure a legacy network adapter to PXE boot a generation 1 VM.*
6. *Windows Server 2012 and Windows Server 2012 R2 support NIC teaming as a feature.*
7. *You must enable MAC address spoofing if you are going to create a NIC team within the VM guest operating system.*
8. *You need to configure a static IP address pool as this will allow VMM to assign IP addresses to VMs running Windows operating systems without requiring a DHCP server.*
9. *You need to install the Routing role service, which is available as part of the Remote Access role.*
10. *You would choose to create Private VLAN networks as PVLAN networks support more than the 12-bit address space limitation imposed by traditional VLANs.*

Chapter 4

1. *The VM needs to be in a shutdown state before you can perform failover.*
2. *You would configure power optimization if you wanted an underutilized Hyper-V server powered down automatically by VMM.*
3. *You need to deploy System Center 2012/2012 R2 Operations Manager before you can use Performance and Resource Optimization.*
4. *The guest OS profile allows you to configure the built-in administrator account password, computer name, as well as roles and features that should be installed.*
5. *You configure a hardware profile to specify the number of virtual processors and virtual network adapters a VM is configured with when deployed.*
6. *You can create a VMM template from an existing virtual hard disk, an existing VMM template stored in the VMM library or from an existing VM deployed on a host.*

Chapter 5

1. *You would select the Fabric Administrator role profile as this allows you to assign permission to perform all administrative tasks over one or more host groups.*
2. *You would select the Read-Only Administrator role profile as this allows users assigned the role to view, but not alter, settings, status and job status information related to VMM objects.*
3. *The Tenant Administrator and Application Administrator role profiles can only be assigned private cloud scopes.*
4. *You should configure a member level quota as this applies quotas to individual role members.*
5. *You will configure a role level quota as this limits combined resource utilization for all role members.*
6. *You need to configure the virtual hard disk as an equivalent object and then configure the template to use this virtual hard disk.*
7. *You must install the Operations Manager console on the VMM management server if you want to integrate Operations Manager with VMM.*

Chapter 6

1. *SAS, iSCSI, Fibre Channel and Fibre Channel over Ethernet.*
2. *You can use NTFS or ReFS with disk witnesses in Windows Server 2012 R2 failover clusters.*
3. *The node majority quorum model is used by default with failover cluster deployments that have an odd number of nodes.*
4. *You would deploy an even number of nodes if you had a disk witness and wanted to deploy a cluster that uses the Node and Disk Majority quorum model.*
5. *Force Quorum Resiliency is a feature of Windows Server 2012 R2 that minimizes the chance that a cluster will be in a partitioned state.*
6. *Remote updating mode allows you to monitor the Cluster Aware Updating process.*
7. *Must use .vhdx format, must be stored on a scale-out file server or Cluster Shared Volume, must be connected to guest VMs through a virtual SCSI controller, guest OS must be Windows Server 2012 or Windows Server 2012 R2, Enable virtual hard disk sharing must be enabled.*

Chapter 7

1. *RDMA (Remote Direct Memory Access) must be supported and enabled to use SMB Direct with VM live migration.*
2. *You must configure constrained delegation on the cifs and Microsoft Virtual System Migration Service services if you want to use Kerberos to authenticate live migration traffic.*
3. *You should select the Copy the virtual machine (create a new unique ID) as this will allow you to have both VMs on the same Hyper-V host.*
4. *You should enable the network protection feature as this will automatically live migrate the VM to another node should connectivity be lost on the current node and available on the destination node.*
5. *The file share must be an SMB 3.0 share to be able to host Hyper-V configuration, checkpoint, virtual hard disk and smart paging files.*
6. *The VM drain on shutdown feature is a Windows Server 2012 R2 feature that ensures that running VMs are live migrated off a cluster node if the node is restarted or shut down without being put into maintenance mode.*
7. *System Center 2012 R2 VMM does not support P2V migrations.*
8. *Prior to performing a V2V migration, you'll need to disable anti-malware software, remove VMware tools and shut down the VM.*

Chapter 7

1. *Sealed management packs use the .mp(b) extension.*
2. *SQL Server Reporting Services must be installed before you can configure and deploy reporting in Operations Manager.*
3. *Audit Collection Services allows you to centralize the collection of auditing data.*
4. *The Fabric Health Dashboard provides you information about the health of each VMM private cloud as well as the fabric that supports those clouds.*
5. *The domain controller should be a member of the Cloneable Domain Controllers group.*
6. *You should use Kerberos authentication if you want data sent over the network in an unencrypted format on port 80.*
7. *You can configure Hyper-V replica to store up to 24 recovery points.*
8. *The VM must be in a shutdown state before you can perform a planned failover to the replica host.*
9. *System Center Global Service Monitor provides monitoring of external web applications from multiple locations around the world.*

About the Author



Orin Thomas was born in 1973. He lives with his wife Oksana and son Rooslan in Melbourne, Australia.

Orin works as an author and has written more than a thirty IT textbooks. He is the convener of the [Melbourne System Center, Security, and Infrastructure Group](#) and a Microsoft Security MVP. He also works as an author for [PluralSight](#).

He is currently working on updating several Windows Server 2012 books to Windows Server 2012 R2

He writes a blog for Windows IT Pro magazine called [Hyperbole, Embellishment, and Systems Administration](#).

Follow him on twitter [@orinthomas](#).

About Veeam Software

Veeam[®] is Protection for the Modern Data Center[™] - providing powerful, easy-to-use and affordable solutions that are Built for Virtualization[™] and the Cloud. [Veeam Backup & Replication[™]](#) delivers [VMware vSphere backup](#), [Hyper-V backup](#), recovery and replication. This #1 VM Backup[™] solution helps organizations meet RPOs and RTOs, save time, eliminate risks and dramatically reduce capital and operational costs. [Veeam Backup Management Suite[™]](#) provides all the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure. [Veeam Management Pack[™]](#) (MP) extends enterprise monitoring to vSphere through Microsoft System Center and also offers monitoring and reporting for the Veeam Backup & Replication infrastructure. The [Veeam Cloud Provider Program](#) (VCP) offers flexible monthly and perpetual licensing to meet the needs of hosting, managed service and cloud service providers. VCP currently has over 4,000 service provider participants worldwide. Monthly rental is available in more than 70 countries from more than 50 Veeam aggregators.

Founded in 2006, Veeam currently has 23,000 ProPartners and more than 91,500 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>.



Protection for the
Modern Data Center



To learn more, visit <http://www.veeam.com/backup>