

Contents

Overview

[What is VMM?](#)

[What's new in VMM](#)

Get Started

[Release notes - VMM](#)

[Turn telemetry data on/off](#)

[Deploy a VMM cloud](#)

[Create a VMM cloud](#)

[Manage a VMM cloud](#)

[Deploy a guarded host fabric](#)

[Deploy guarded hosts](#)

[Configure fallback HGS settings](#)

[Deploy a shielded VHDX and VM template](#)

[Deploy a shielded VM](#)

[Deploy a shielded Linux VM](#)

[Deploy and manage a software defined network \(SDN\) infrastructure](#)

[Deploy an SDN network controller](#)

[Deploy an SDN SLB](#)

[Deploy an SDN RAS gateway](#)

[Deploy SDN using PowerShell](#)

[Set up a VM network in SDN](#)

[Encrypt VM networks in SDN](#)

[Allow and block VM traffic with SDN port ACLs](#)

[Control SDN virtual network bandwidth with QoS](#)

[Load balance network traffic](#)

[Set up NAT for traffic forwarding in an SDN](#)

[Route traffic across networks in the SDN infrastructure](#)

[Configure SDN guest clusters](#)

[Update the NC server certificate](#)

- Set up SDN SLB VIPs

- Back up and restore the SDN infrastructure

- Remove an SDN from VMM

- Manage SDN resources in the VMM fabric

- Deploy and manage Storage Spaces Direct

- Set up a hyper-converged Storage Spaces Direct cluster

- Set up a disaggregated Storage Spaces Direct cluster

- Manage Storage Spaces Direct clusters

- Assign storage QoS policies for Clusters

- How To

- Plan

- System requirements - VMM

- Plan VMM installation

- Plan a VMM high availability deployment

- Identify VMM ports and protocols

- Plan the VMM compute fabric

- Plan the VMM networking fabric

- Identify supported storage arrays

- Upgrade and install

- Upgrade VMM

- Install VMM

- Install the VMM console

- Enable enhanced console session

- Deploy VMM for high availability

- Deploy a highly available VMM management server

- Deploy a highly available SQL Server database for VMM

- Deploy a highly available VMM library

- Set up TLS 1.2

- Deploy update rollups

- Back up and restore VMM

- Manage the VMM library

- Library overview

Add file-based resources to the VMM library

Add profiles to the VMM library

Add VM templates to the VMM library

Add service templates to the VMM library

Manage VMM library resources

Manage virtualization servers

Manage VMM host groups

Add existing Hyper-V hosts and clusters to the fabric

Add a Nano server as a Hyper-V host or cluster

Run a script on host

Create a cluster from standalone Hyper-V hosts

Provision a Hyper-V host or cluster from bare-metal

Create a guest Hyper-V cluster from a service template

Set up networking for Hyper-V hosts and clusters

Set up storage for Hyper-V hosts and clusters

Manage MPIO for Hyper-V hosts and clusters

Manage Hyper-V extended port ACLs

Manage Hyper-V clusters

Update Hyper-V hosts and clusters

Run a rolling upgrade of Hyper-V clusters

Service Hyper-V hosts for maintenance

Manage VMware servers

Manage management servers

Manage infrastructure servers

Manage update servers

Manage networking

Network fabric overview

Set up logical networks

Set up VM networks

Set up IP address pools

Add a network gateway

Set up port profiles

Set up logical switches

Set up MAC address pools

Integrate NLB with service templates

Set up an IPAM server

Manage storage

Set up storage fabric

Set up storage classifications

Add storage devices

Allocate storage to host groups

Set up a Microsoft iSCSI Target Server

Set up a Virtual Fibre Channel

Set up file storage

Set up Storage Replica in VMM

Manage SOFS

Set up a scale-out file server (SOFS)

Perform a rolling upgrade on SOFS

Add an existing SOFS to the VMM fabric

Create an SOFS cluster from standalone servers

Provision SOFS from bare-metal computers

Manage SOFS settings

Manage virtual machines

Provision VMs

Deploy VMs from a blank VHD

Deploy VMs from an existing VHD

Clone existing VMs

Deploy VMs with rapid provisioning using SAN copy

Deploy VMs from a VM template

Deploy Linux VMs

Deploy nested VMs

Convert VMware VMs to Hyper-V

Install an operating system on a VM

Manage VM settings

- Manage dynamic and power optimization for VMs

- Create VM role templates with VMM and Azure Pack

Migrate VMs

- Migration overview

- Migrate a VM

- Migrate VM storage

- Run a live migration

Manage roles and accounts

- Roles and accounts overview

- Set up user roles

- Set up Run As accounts

- Manage self-service users

- Perform self-service user tasks

Monitor VMM

- Monitoring overview

- Set up monitoring and reporting with Operations Manager

Integrate with Azure

- Add an Azure subscription in VMM

- VM update management

- Manage Azure VMs

- Manage Azure Resource Manager-based and region-specific Azure subscriptions

What is Virtual Machine Manager?

3/4/2019 • 2 minutes to read

Welcome to System Center Virtual Machine Manager (VMM)! VMM is part of the System Center suite, used to configure, manage and transform traditional datacenters, and helping to provide a unified management experience across on-premises, service provider, and the Azure cloud. VMM capabilities include:

- **Datacenter:** Configure and manage your datacenter components as a single fabric in VMM. Datacenter components include virtualization servers, networking components, and storage resources. VMM provisions and manages the resources needed to create and deploy virtual machines and services to private clouds.
- **Virtualization hosts:** VMM can add, provision, and manage Hyper-V and VMware virtualization hosts and clusters.
- **Networking:** Add networking resources to the VMM fabric, including network sites defined by IP subnets, virtual LANs (VLANs), logical switches, static IP address and MAC pools. VMM provides network virtualization, including support for creating and manage virtual networks and network gateways. Network virtualization allows multiple tenants to have isolated networks and their own IP address ranges for increased privacy and security. Using gateways, VMs on virtual networks can connect to physical networks in the same site or in different locations.
- **Storage:** VMM can discover, classify, provision, allocate, and assign local and remote storage. VMM supports block storage (fibre channel, iSCSI, and Serial Attached SCSI (SAS) storage area networks (SANs)).
- **Library resources:** The VMM fabric retains a library of file-based and non file-based resources that are used to create and deploy VMs and services on virtualization hosts. File-based resources include virtual hard disks, ISO images, and scripts. Non file-based resources include templates and profiles that are used to standardize the creation of VMs. Library resources are accessed through library shares.

Resources

- For VMM questions or comments, go to [System Center Virtual Machine Manager Forums](#).
- To read blog posts from the VMM engineering team, see [System Center: Virtual Machine Manager Engineering Blog](#).

Next steps

- [Learn about](#) system requirements.

What's new in System Center Virtual Machine Manager

3/14/2019 • 23 minutes to read

This article details the new features supported in System Center 2019 - Virtual Machine Manager (VMM).

New features in VMM 2019

See the following sections for detailed information about the new features/feature updates supported in VMM 2019.

Compute

Cluster Rolling upgrade for S2D clusters

System Center 2019 - Virtual Machine Manager (VMM) supports rolling upgrade of a Storage Spaces Direct (S2D) host cluster from Windows Server 2016 to Windows Server 2019. [Learn More](#).

Support for deduplication for ReFS volume

VMM 2019 supports deduplication for ReFS volume on WS 2019 hyper-converged cluster and SOFS. [Learn more](#).

Storage

Storage DO

This feature helps in preventing cluster shared storage (CSV and file shares) from becoming full due to expansion/new VHDs placed on the cluster shared storage. You can now set a threshold value to trigger a warning when free storage space in the cluster shared storage falls below the threshold, during a new disk placement or auto migration of VHDs to other shared storage in the cluster. [Learn more](#).

Support for storage health monitoring

Storage health monitoring helps you to monitor the health and operational status of storage pool, LUNs, and physical disks in the VMM fabric.

You can monitor the storage health in the **Fabric** page of VMM console. [Learn more](#).

Networking

Configuration of SLB VIPs through VMM service templates

SDN in Windows 2016 can use Software Load Balancing (SLB) to evenly distribute network traffic among workloads managed by service provider and tenants. VMM 2016 currently supports deployment of SLB Virtual IPs (VIPs) using power shell.

With VMM 2019, VMM supports configuration of SLB VIPs while deploying multi-tier application by using the service templates. [Learn more](#).

Configuration of encrypted VM networks through VMM

VMM 2019 supports encryption of VM networks. Using the new encrypted networks feature, end-to-end encryption can be easily configured on VM networks by using the network controller (NC). This encryption prevents the traffic between two VMs on the same network and same subnet, from being read and manipulated.

The control of encryption is at the subnet level and encryption can be enabled/disabled for each subnet of the VM network. [Learn more.](#)

Support for configuring layer 3 forwarding gateway using VMM console

L3 forwarding enables connectivity between the physical infrastructure in the datacenter and the virtualized infrastructure in the Hyper-V network virtualization cloud. Earlier versions of VMM supported the Layer 3 gateway configuration through PowerShell.

In VMM 2019, you can configure Layer 3 forwarding gateway using the VMM console. [Learn more.](#)

Support for Static MAC address on VMs deployed on a VMM cloud

This feature allows you to set static MAC address on VMs deployed on a cloud. You can also change the MAC address from static to dynamic and vice versa for the already deployed VMs. [Learn more.](#)

Azure Integration

VM update management through VMM using Azure Automation Subscription

VMM 2019 is introducing the possibility of patching and updating on-prem VMs (managed by VMM) by integrating VMM with Azure automation subscription. [Learn more.](#)

New RBAC Role - Virtual Machine Administrator

In a scenario where enterprises want to create a user role for troubleshooting, it is necessary that the user has access to all the VMs so the user can make any required changes on the VMs to resolve the issue. There is also a need for the user to have access to the fabric to identify the root cause for the issue. However, for security reasons, this user should not be given the privileges to make any changes on the fabric (such as add storage, add hosts etc).

The current role-based access control (RBAC) in VMM does not have a role defined for this persona and the existing roles of Delegated Admin and Fabric admin have too little or more than necessary permissions to perform just troubleshooting.

To address this issue, VMM 2019 supports a new role called **Virtual Machine Administrator**. The user of this role has Read and Write access to all VMs but read-only access to the fabric. [Learn more.](#)

Support for Group Managed Service Account as VMM service account

Group Managed Service Account (gMSA) helps improve the security posture and provides convenience through automatic password management, simplified service principle name (SPN) management, and the ability to delegate the management to other administrators.

VMM 2019 supports the use of gMSA for **Management server service account**. [Learn more.](#)

NOTE

The following features/feature updates were introduced in VMM 1807, included in VMM 2019.

Storage

Supports selection of CSV for placing a new VHD

VMM allows you to select a cluster shared volumes (CSV) for placing a new virtual hard disc (VHD).

In earlier versions of VMM, a new VHD on a virtual machine (VM), by default, is placed on the same CSV where the earlier VHDs associated with the VM are placed, there was no option to choose a different CSV/ folder. In case of any issues related to the CSV, such as storage full or over commitment, users had to migrate the VHD, only after deploying the VHD.

With VMM 1807, you can now choose any location to place the new disc. You can manage this disc easily, based on

the storage availability of CSVs. [Learn more.](#)

Networking

Display of LLDP information for networking devices

VMM supports Link Layer Discovery Protocol (LLDP). You can now view network device properties and capabilities information of the hosts from VMM. Host operating system must be Windows 2016 or higher.

DataCenterBridging and DataCenterBridging-LLDP-Tools features have been enabled on hosts to fetch the LLDP properties. [Learn more.](#)

Convert SET switch to logical switch

VMM allows you to convert a switch embedded teaming (SET) switch to logical switch by using the VMM console. In earlier versions, this feature was supported only through PowerShell script. [Learn more.](#)

VMware host management

VMM supports VMware ESXi v6.5 servers in VMM fabric. This support facilitates the administrators with additional flexibility in managing multiple hypervisors in use. [Learn more](#) about the additional details of supported vmware server versions.

Support for S2D cluster update

VMM supports update of an S2D host or a cluster. You can update individual S2D hosts or clusters against the baselines configured in windows server update services (WSUS). [Learn more.](#)

Others

Support for SQL 2017

VMM supports SQL 2017. You can upgrade SQL 2016 to 2017.

NOTE

The following features/feature updates were introduced in VMM 1801, included in VMM 2019.

Compute

Nested virtualization

VMM supports Nested Virtualization feature that allows you to run Hyper-V inside a Hyper-V virtual machine. In other words, with nested virtualization, a Hyper-V host itself can be virtualized. Nested virtualization can be enabled out-of-band by using PowerShell and Hyper-V host configuration.

You can leverage this functionality to reduce your infrastructure expense for development, test, demo, and training scenarios. This feature also allows you to use third-party virtualization management products with Microsoft hypervisor.

You can enable or disable the nested virtualization feature using VMM. You can configure the VM as a Host in VMM and perform host operations from VMM on this VM. For example, VMM dynamic optimization considers a nested VM host for placement. [Learn more.](#)

Migration of VMware VM (EFI firmware-based VM) to Hyper-V VM

The current VMM migration for VMware VMs to Hyper-V only supports migration of BIOS-based VMs.

VMM enables migration of EFI based VMware VMs to Hyper-V generation 2 VMs. VMware VMs that you migrate to Microsoft Hyper-V platform can take advantage of the Hyper-V generation 2 features.

As part of this release, the **Convert Virtual machine** wizard enables the VM migration based on the firmware type (BIOS or EFI), selects and defaults the Hyper-V VM generation appropriately: [Learn more](#).

1. BIOS-based VMs are migrated to Hyper-V VM generation 1.
2. EFI-based VMs are migrated to Hyper-V VM generation 2.

We have also made improvements in the VMWare VM conversion process that makes the conversion up to 50% faster.

Performance improvement in host refresher

VMM host refresher has undergone certain updates for performance improvement.

With these updates, in scenarios where the organization is managing large number of hosts and VMs with checkpoints – you would be able to observe significant and noticeable improvements in the performance of the job.

In our lab with VMM instances managing 20 hosts - each host managing 45-100 VMs, we have measured up to 10X performance improvement.

Enhanced console session in VMM

Console connect in VMM provides an alternative way to remote desktop to connect to the VM. This is most useful when the VM does not have any network connectivity or want to change network configuration that could break the network connectivity. Currently, the current console connect in VMM supports only basic session where clipboard text can only be pasted through Type Clipboard Text menu option.

VMM supports enhanced console session that enables **Cut (Ctrl + X)**, **Copy (Ctrl + C)** and **Paste (Ctrl + V)** operations on the ANSI text and files available on the clipboard, thereby copy/paste commands for text and files are possible from and to the VM. [Learn more](#).

Storage

Improvement in VMM storage QoS

Storage Quality of Service (SQoS) provides a way to centrally monitor and manage storage performance for virtual machines using Hyper-V and the Scale-Out File Server (SOFS) roles. The feature automatically improves storage resource fairness between multiple VMs using the same cluster and allows policy-based performance goals.

VMM supports the following improvements in SQoS:

- Extension of SQoS support beyond S2D - You can now assign storage QoS policies to storage area networks (SAN). [Learn more](#).
- Support for VMM private cloud - storage QoS policies can now be consumed by the VMM cloud tenants. [Learn more](#).
- Availability of storage QoS policies as templates - You can set storage QoS policies through VM templates. [Learn more](#).

Networking

Configuration of guest clusters in SDN through VMM

With the advent of the software defined network (SDN), in Windows Server 2016 and System Center 2016, the configuration of guest clusters has undergone some change.

With the introduction of the SDN, VMs which are connected to the virtual network using SDN are only permitted to use the IP address that the network controller assigns for communication. The SDN design is inspired by Azure networking design, supports the floating IP functionality through the Software Load Balancer (SLB), like Azure

networking.

VMM also supports the floating IP functionality through the Software Load Balancer (SLB) in the SDN scenarios. VMM 1801 supports guest clustering through an Internal Load Balancer (ILB) VIP. The ILB uses probe ports, which are created on the guest cluster VMs to identify the active node. At any given time, the probe port of only the active node responds to the ILB, and all the traffic directed to the VIP is routed to the active node. . [Learn more](#).

Configuration of SLB VIPs through VMM service templates

SDN in Windows 2016 can use Software Load Balancing (SLB) to evenly distribute network traffic among workloads managed by service provider and tenants. VMM 2016 currently supports deployment of SLB Virtual IPs (VIPs) using power shell.

VMM supports configuration of SLB VIPs while deploying multi-tier application by using the service templates. [Learn more](#).

Configuration of encrypted VM networks through VMM

VMM supports encryption of VM networks. Using the new encrypted networks feature, end-to-end encryption can be easily configured on VM networks by using the Network Controller (NC). This encryption prevents traffic between two VMs on the same network and same subnet, from being read and manipulated.

The control of encryption is at the subnet level and encryption can be enabled/disabled for each subnet of the VM network. [Learn more](#).

Security

Support to Linux shielded VM

Windows Server 2016 introduced the concept of a shielded VM for Windows OS-based VMs. Shielded VMs provide protection against malicious administrator actions both when VM's data is at rest or an untrusted software is running on Hyper-V hosts.

With Windows Server 1709, Hyper-V introduces support for provisioning Linux shielded VMs and the same has been extended to VMM. [Learn more](#).

Configuration of fallback HGS

Being at the heart of providing attestation and key protection services to run shielded VMs on Hyper-V hosts, the host guardian service (HGS) should operate even in situations of disaster. Windows Server 1709 added support for fallback HGS.

Using VMM, a guarded host can be configured with a primary and a secondary pair of HGS URLs (an attestation and key protection URI). This capability enables scenarios such as guarded fabric deployments spanning two data centers for disaster recovery purposes, HGS running as shielded VMs etc.

The primary HGS URLs will always be used in favor of the secondary. If the primary HGS fails to respond after the appropriate timeout and retry count, the operation will be reattempted against the secondary. Subsequent operations will always favor the primary; the secondary will only be used when the primary fails. [Learn more](#).

Azure Integration

Management of Azure Resource Manager-based and region-specific Azure subscriptions

Currently, the VMM Azure plugin supports only classic virtual machines (VMs) and public Azure regions.

VMM 1801 supports management of Azure Resource Manager based VMs, Azure Active Directory (AD) based authentication that is created by using the new Azure portal and region-specific Azure subscriptions (namely, Germany, China, US Government Azure regions). [Learn more](#).

This article details the new features supported in System Center 1807 - Virtual Machine Manager (VMM).

What's new in System Center 1807 - Virtual Machine Manager

See the following sections for information about the new features supported in VMM 1807.

NOTE

To view the bugs fixed and the installation instructions for VMM 1807, see [KB article 4135364](#).

Storage

Supports selection of CSV for placing a new VHD

VMM 1807 allows you to select a cluster shared volumes (CSV) for placing a new virtual hard disc (VHD).

In earlier versions of VMM, a new VHD on a virtual machine (VM), by default, is placed on the same CSV where the earlier VHDs associated with the VM are placed, there was no option to choose a different CSV/ folder. In case of any issues related to the CSV, such as storage full or over commitment, users had to migrate the VHD, only after deploying the VHD.

With VMM 1807, you can now choose any location to place the new disc. You can manage this disc easily, based on the storage availability of CSVs. [Learn more](#).

Networking

Display of LLDP information for networking devices

VMM 1807 supports Link Layer Discovery Protocol (LLDP). You can now view network device properties and capabilities information of the hosts from VMM. Host operating system must be Windows 2016 or higher.

DataCenterBridging and DataCenterBridging-LLDP-Tools features have been enabled on hosts to fetch the LLDP properties. [Learn more](#).

Convert SET switch to logical switch

VMM 1807 allows you to convert a switch embedded teaming (SET) switch to logical switch by using the VMM console. In earlier versions, this feature was supported only through PowerShell script. [Learn more](#).

VMware host management

VMM 1807 supports VMware ESXi v6.5 servers in VMM fabric. This support facilitates the administrators with additional flexibility in managing multiple hypervisors in use. [Learn more](#) about the additional details of supported vmware server versions.

Support for S2D cluster update

VMM 1807 supports update of an S2D host or a cluster. You can update individual S2D hosts or clusters against the baselines configured in windows server update services (WSUS). [Learn more](#).

Others

Support for SQL 2017

VMM 1807 supports SQL 2017. You can upgrade SQL 2016 to 2017.

This article details the new features supported in System Center 1801 - Virtual Machine Manager (VMM).

This article details the new features supported in System Center 2016 - Virtual Machine Manager (VMM).

What's new in System Center 1801 - Virtual Machine Manager

See the following sections for detailed information about the new features supported in VMM 1801.

Compute

Nested virtualization

VMM supports Nested Virtualization feature that allows you to run Hyper-V inside a Hyper-V virtual machine. In other words, with nested virtualization, a Hyper-V host itself can be virtualized. Nested virtualization can be enabled out-of-band by using PowerShell and Hyper-V host configuration.

You can leverage this functionality to reduce your infrastructure expense for development, test, demo, and training scenarios. This feature also allows you to use third- party virtualization management products with Microsoft hypervisor.

You can enable or disable the nested virtualization feature using SCVMM 1801. You can configure the VM as a Host in VMM and perform host operations from VMM on this VM. For example, VMM dynamic optimization considers a nested VM host for placement. [Learn more](#).

Migration of VMware VM (EFI firmware-based VM) to Hyper-V VM

The current VMM migration for VMware VMs to Hyper-V only supports migration of BIOS-based VMs.

VMM 1801 release enables migration of EFI based VMware VMs to Hyper-V generation 2 VMs. VMware VMs that you migrate to Microsoft Hyper-V platform can take advantage of the Hyper-V generation 2 features.

As part of this release, the **Convert Virtual machine** wizard enables the VM migration based on the firmware type (BIOS or EFI), selects and defaults the Hyper-V VM generation appropriately: [Learn more](#).

1. BIOS-based VMs are migrated to Hyper-V VM generation 1.
2. EFI-based VMs are migrated to Hyper-V VM generation 2.

We have also made improvements in the VMWare VM conversion process that makes the conversion up to 50% faster.

Performance improvement in host refresher

The VMM 1801 host refresher has undergone certain updates for performance improvement.

With these updates, in scenarios where the organization is managing large number of hosts and VMs with checkpoints – you would be able to observe significant and noticeable improvements in the performance of the job.

In our lab with VMM instances managing 20 hosts - each host managing 45-100 VMs, we have measured up to 10X performance improvement.

Enhanced console session in VMM

Console connect in VMM provides an alternative way to remote desktop to connect to the VM. This is most useful when the VM does not have any network connectivity or want to change network configuration that could break the network connectivity. Currently, the current console connect in VMM supports only basic session where clipboard text can only be pasted through Type Clipboard Text menu option.

VMM 1801 supports enhanced console session that enables **Cut (Ctrl + X)**, **Copy (Ctrl + C)** and **Paste (Ctrl + V)** operations on the ANSI text and files available on the clipboard, thereby copy/paste commands for text and files are possible from and to the VM. [Learn more](#).

Storage

Improvement in VMM storage QoS

Storage Quality of Service (SQoS) provides a way to centrally monitor and manage storage performance for

virtual machines using Hyper-V and the Scale-Out File Server (SOFS) roles. The feature automatically improves storage resource fairness between multiple VMs using the same cluster and allows policy-based performance goals.

VMM 1801 supports the following improvements in SQoS:

- Extension of SQoS support beyond S2D - You can now assign storage QoS policies to storage area networks (SAN). [Learn more.](#)
- Support for VMM private cloud - storage QoS policies can now be consumed by the VMM cloud tenants. [Learn more.](#)
- Availability of storage QoS policies as templates - You can set storage QoS policies through VM templates. [Learn more.](#)

Networking

Configuration of guest clusters in SDN through VMM

With the advent of the software defined network (SDN), in Windows Server 2016 and System Center 2016, the configuration of guest clusters has undergone some change.

With the introduction of the SDN, VMs which are connected to the virtual network using SDN are only permitted to use the IP address that the network controller assigns for communication. The SDN design is inspired by Azure networking design, supports the floating IP functionality through the Software Load Balancer (SLB), like Azure networking.

VMM 1801 release also supports the floating IP functionality through the Software Load Balancer (SLB) in the SDN scenarios. VMM 1801 supports guest clustering through an Internal Load Balancer (ILB) VIP. The ILB uses probe ports, which are created on the guest cluster VMs to identify the active node. At any given time, the probe port of only the active node responds to the ILB and all the traffic directed to the VIP is routed to the active node. . [Learn more.](#)

Configuration of SLB VIPs through VMM service templates

SDN in Windows 2016 can use Software Load Balancing (SLB) to evenly distribute network traffic among workloads managed by service provider and tenants. VMM 2016 currently supports deployment of SLB Virtual IPs (VIPs) using power shell.

With VMM 1801, VMM supports configuration of SLB VIPs while deploying multi-tier application by using the service templates. [Learn more.](#)

Configuration of encrypted VM networks through VMM

VMM 1801 supports encryption of VM networks. Using the new encrypted networks feature, end-to-end encryption can be easily configured on VM networks by using the Network Controller (NC). This encryption prevents traffic between two VMs on the same network and same subnet, from being read and manipulated.

The control of encryption is at the subnet level and encryption can be enabled/disabled for each subnet of the VM network. [Learn more.](#)

Security

Support to Linux shielded VM

Windows Server 2016 introduced the concept of a shielded VM for Windows OS-based VMs. Shielded VMs provide protection against malicious administrator actions both when VM's data is at rest or an untrusted software is running on Hyper-V hosts.

With Windows Server 1709, Hyper-V introduces support for provisioning Linux shielded VMs and the same has been extended to VMM 1801. [Learn more.](#)

Configuration of fallback HGS

Being at the heart of providing attestation and key protection services to run shielded VMs on Hyper-V hosts, the host guardian service (HGS) should operate even in situations of disaster. Windows Server 1709 added support for fallback HGS.

Using VMM 1801, a guarded host can be configured with a primary and a secondary pair of HGS URLs (an attestation and key protection URI). This capability enables scenarios such as guarded fabric deployments spanning two data centers for disaster recovery purposes, HGS running as shielded VMs etc.

The primary HGS URLs will always be used in favor of the secondary. If the primary HGS fails to respond after the appropriate timeout and retry count, the operation will be reattempted against the secondary. Subsequent operations will always favor the primary; the secondary will only be used when the primary fails. [Learn more.](#)

Azure Integration

Management of Azure Resource Manager-based and region-specific Azure subscriptions

Currently, the VMM Azure plugin supports only classic virtual machines (VMs) and public Azure regions.

VMM 1801 supports management of Azure Resource Manager based VMs, Azure Active Directory (AD) based authentication that is created by using the new Azure portal and region-specific Azure subscriptions (namely, Germany, China, US Government Azure regions). [Learn more.](#)

What's new in VMM 2016

See the following sections for detailed information about the new features supported in VMM 2016.

Compute

Full lifecycle management of Nano Server-based hosts and VMs

You can provision and manage Nano Server-based hosts and virtual machines in the VMM fabric. [Learn more.](#)

Rolling upgrade of a Windows Server 2012 R2 host clusters

You can now upgrade Hyper-V and scale-out file server (SOFS) clusters in the VMM fabric from Windows Server 2012 R2 to Windows Server 2016, with no downtime for the host workloads. VMM orchestrates the entire workflow. It drains the node, removes it from the cluster, reinstalls the operating system, and adds it back into the cluster. Learn more about performing rolling upgrades for [Hyper-V clusters](#), and [SOFS clusters](#).

Creating Hyper-V & SOFS clusters

There's a streamlined workflow for creating Hyper-V and SOFS clusters:

- **Bare metal deployment of Hyper-V host clusters:** Deploying a Hyper-V host cluster from bare metal machines is now a single step. [Learn more](#)
- **Adding a bare-metal node to an existing Hyper-V host cluster or an SOFS Cluster:** You can now directly add a bare-metal computer to an existing [Hyper-V](#) or [SOFS](#) cluster.

New operations for running VMs

You can now increase/decrease static memory and add/remove virtual network adapter for virtual machines that are running. [Learn more.](#)

Production checkpoints

You can now create production checkpoints for VMs. These checkpoints are based on Volume Shadow Copy Service (VSS) and are application-consistent (compared to standard checkpoints based on saved state technology that aren't. [Learn more.](#)

Server App-V

The Server App-V application in service templates is no longer available in VMM 2016. You can't create new templates or deploy new services with the Server App-V app. If you upgrade from VMM 2012 R2 and have a service with the Server App-V application, the existing deployment will continue to work. However, after the upgrade you can't scale out the tier with Server App-V application. You can scale out other tiers.

Storage

Deploy and manage storage clusters with Storage Spaces Direct (S2D)

[Storage Spaces Direct in Windows Server 2016](#) enables you to build highly available storage systems on Windows Server. You can use VMM to create a Scale-Out File Server running Windows Server 2016, and configure it with Storage Spaces Direct. After it's configured you can create storage pools and file shares on it. [Learn more](#).

Storage Replica

In VMM 2016 you can use Windows Storage Replica to protect data in a volume by synchronously replicating it between primary and secondary (recovery) volumes. You can deploy the primary and secondary volumes to a single cluster, to two different clusters, or to two standalone servers. You use PowerShell to set up Storage Replica and run failover. [Learn more](#)

Storage Quality of Service (QoS)

You can configure QoS for storage to ensure that disks, VMs, apps, and tenants don't drop below a certain resource quality when hosts and storage are handling heavy loads. You can configure QoS for storage in the VMM fabric.

Networking

Software Defined Networking (SDN)

In VMM 2016 you can deploy the entire SDN stack using VMM service templates.

- You can deploy and manage a multi-node Network Controller in a subnet. After you deploy and onboard the Network Controller, you can specify that fabric components should be managed with SDN, to provide connectivity to tenant VMs and to define policies.
- You can deploy and configure a software load balancer, to distribute traffic within networks managed by Network Controller. The software load balancer can be used for inbound and outbound NAT.
- You can deploy and configure a Windows Server Gateway pool with M+N redundancy. After you deploy the gateway, you connect a tenant network to a hosting provider network, or to your own remote data center network using S2S GRE, S2S IPsec, or L3.

Network traffic isolation and filtering

You can limit and segregate network traffic by specifying port ACLs on VM networks, virtual subnets, network interfaces, or on an entire VMM stamp using Network Controller and PowerShell. [Learn more](#).

Virtual network adapter naming

When you deploy a virtual machine, you might want to run a post-deployment script on the guest operating system to configure virtual network adapters. Previously, this was difficult because there wasn't an easy way to distinguish different virtual network adapters during deployment. Now, for generation 2 virtual machines deployed on Hyper-V hosts running Windows Server 2016, you can name the virtual network adapter in a virtual machine template. This is similar to using consistent device naming (CDN) for a physical network adapter.

Self-service SDN management using Windows Azure Pack (WAP)

You can provide self-service capabilities for fabric managed by Network Controller. These include creating and managing VM networks, configuring S2S IPsec connections, and configuring NAT options for tenant and infrastructure VMs in your data center.

Logical switch deployment across hosts

- The interface for creating a logical switch has been streamlined to make it easier to select settings.
- You can directly use Hyper-v to configure a standard virtual switch on a managed host, and then use VMM to convert the standard virtual switch to a VMM logical switch, which you later apply on additional hosts.
- When apply a logical switch to a particular host, if the entire operation doesn't succeed, the operation is reverted and host settings are left unchanged. Improved logging makes it easier to diagnose failures.

Security

Guarded host deployment

You can provision and manage guarded hosts and shielded VMs in the VMM fabric, to help provide protection against malicious host administrators and malware.

- You can manage guarded hosts in the VMM compute fabric. You configure guarded hosts to communicate with HGS servers, and you can specify code integrity policies that restrict software that can run in kernel mode on the host.
- You can convert existing VMs to shielded VMs, and deploy new shielded VMs.

Next steps

- [Know the VMM system requirements](#)
- [Plan VMM installation](#)

Release notes for System Center Virtual Machine Manager

3/14/2019 • 17 minutes to read

This article lists the release notes for System Center 2019 - Virtual Machine Manager (VMM).

VMM 2019 release notes

The following sections summarize the release notes for VMM 2019 and include the known issues and workarounds.

Removal of cluster node fails with CleanUpDisks flag

Description: When you remove a cluster node from Windows Server 2019 S2D cluster, with a CleanUpDisks flag, the removal fails with **Could not get the specified instance MSFT_StorageJob** error, in the following scenarios

- The storage capacity is inadequate in the remaining servers to accommodate all the volumes.
- There are not enough fault domains to provide the resiliency of the volume.

Workaround: Ensure the following:

- Adequate storage capacity is available in the remaining servers to accommodate all the volumes
- Enough fault domains are available to provide the resiliency of your volumes.

Addition of storage device having SMI-S management interface fails

Description: Addition of storage device having SMI-S management interface fails with the error *Registration of storage provider failed with error code WsManMIIInvokeFailed* when System Center Virtual Machine Manager (VMM) 2019 is installed on Windows Server 2019.

Workaround: VMM depends on the *Windows Standards-Based Storage Management* service to manage the storage devices using SMI-S. Make sure that the service is started before trying to add the storage device.

Windows Server 2019 does not support HNVv1 networks

Description: Windows Server 2019 does not support HNVv1. If HNVv1 is currently in use, then the cluster that is utilizing HNVv1 should not be upgraded to Windows Server 2019 using Cluster Rolling Upgrade.

Workaround: Migrate out of HNVv1 to SDNv2 on Windows Server 2016 before using Cluster Rolling upgrade to Windows Server 2019.

Latest accessibility fixes in Console are not available

Description: Latest accessibility fixes in Console might not be available when you use .NET 4.7 while installing the VMM console.

Workaround: We recommend you to use .NET 4.8. For detailed information on .NET 4.8 migration, see [the article on .NET migration](#).

Backend adapter connectivity for SLB MUX doesn't work as expected

Description: Backend adapter connectivity of SLB MUX might not work as expected after the migration of virtual machine (VM).

Workaround: Users scale in/scale out in the SLB MUX VM as a workaround.

Cluster Rollup Upgrade fails

Description: Cluster rollup upgrade (CRU) fails during the *Connect Hyper-V host to storage arrays* stage, if the Windows Server 2019 Virtual Hard Disk (VHD) in library server, used as the computer profile for redeploying the operating system (OS) is not installed with the latest updates.

Workaround: To resolve this error, install all the pending updates on the VHD and restart the CRU job.

To avoid this issue, prior to CRU triggering, ensure to install the latest OS updates on the VHD that you want to use for CRU.

Storage Dynamic Optimization does not trigger VHD migration even when optimization criteria are met

Description: Storage Dynamic Optimization (DO) should trigger the VHD migration between Clustered Shared Volumes (CSV), when the free storage space in one of the CSVs falls below the disk space threshold set in the Dynamic Optimization page, and the aggressiveness criteria are met. However, in some cases, the VHDs might not be migrated even if all other Storage DO conditions are met.

Workaround: To ensure Storage migration is triggered, do the following:

1. Check the *HostVolumeID* using *Get-SCStorageVolume* cmdlet. If the *HostVolumeID* returns Null for the volume, refresh the VM and perform Storage DO again.
2. Check the *DiskSpacePlacementLevel* of the host group using the *Get-SCHostReserver* cmdlet. Set the *DiskSpacePlacementLevel* value that is equal to the value of disk space as in *Host Reserve* settings, in the Dynamic Optimization wizard.

Storage Dynamic Optimization disk performs multiple back and forth VHD migrations

Description: If there is a mismatch of disk space warning levels between host groups having the same file share, it can result in multiple migrations, to and from that file share, and might impact storage DO performance.

Workaround: We recommend that you do not do a file share across different clusters, where storage dynamic optimization is enabled.

Performance monitoring for VMM server fails with *Access denied* event error

Description: In a scenario where VMM is monitored using Operations Manager, performance monitoring for VMM server fails with *Access denied* event error. Service users do not have permission to access VirtualMachineManager-Server/Operational event log.

Workaround: Change the security descriptor for Operational event log registry with the following command, and then restart the event log service and health log service.

```
reg add HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\WINEVT\Channels\Microsoft-VirtualMachineManager-Server\Operational /v ChannelAccess /t REG_SZ /d O:BAG:SYD:(D;;0xf0007;;;AN)(D;;0xf0007;;;BG)(A;;0xf0007;;;SY)(A;;0x7;;;BA)(A;;0x3;;;NS)(A;;0x1;;;IU)(A;;0x1;;;SU)"
```

This command will add the service user to the list of allowed users, who can access VirtualMachineManager-Server/Operational event log.

This article lists the release notes for System Center 1807 - Virtual Machine Manager (VMM).

VMM 1807 release notes

The following sections summarize the release notes for VMM 1807 and include the known issues and workarounds.

Latest accessibility fixes in Console are not available

Description: Latest accessibility fixes in Console might not be available when you use .NET 4.7 while installing the VMM Console.

Workaround: We recommend using .NET 4.7.1 while installing the VMM Console. For detailed information on .NET 4.7.1 migration, see [the article on .NET migration](#).

Backend adapter connectivity for SLB MUX doesn't work as expected

Description: Backend adapter connectivity of SLB MUX might not work as expected after VM migration.

Workaround: Users scale in/scale out in the SLB MUX VM as a workaround.

Connectivity issues for SLB addresses

Description: For frontend and backend IP addresses assigned to Software Load Balancer MUX VMs, you might experience connectivity issues if **Register this connection's address in DNS** is selected.

Workaround: Clear the setting to avoid issues with these IP addresses.

VMM integrated with Azure Site Recovery will not support DRA versions earlier than 5.1.3100

Description: In case you are using a VMM integrated with Azure Site Recovery, VMM supports Data Recovery Agent (DRA) version [5.1.3100](#) or higher. Earlier versions are not supported.

Workaround: Use the following steps and upgrade the DRA version:

1. Uninstall existing version of DRA
2. Install VMM 1807 patch)
3. [Install 5.1.3100](#) version or higher.

Host/Cluster refresh might take longer if there are large number of logical network definitions

Description: When there are large number of logical network definitions in the environment, cluster/host refresh might take longer than expected.

This article lists the release notes for System Center 2016 - Virtual Machine Manager (VMM).

VMM 1801 release notes

The following sections summarize the release notes for VMM 1801 and includes the known issues and workarounds.

Latest accessibility fixes in Console are not available

Description: Latest accessibility fixes in Console might not be available when you use .NET 4.7 while installing the VMM Console.

Workaround: We recommend using .NET 4.7.1 while installing the VMM Console. For detailed information on .NET 4.7.1 migration, see [the article on .NET migration](#).

Backend adapter connectivity for SLB MUX doesn't work as expected

Description: Backend adapter connectivity of SLB MUX might not work as expected after VM migration.

Workaround: Users scale in/scale out in the SLB MUX VM as a workaround.

Connectivity issues for SLB addresses

Description: For frontend and backend IP addresses assigned to Software Load Balancer MUX VMs, you might experience connectivity issues if **Register this connection's address in DNS** is selected.

Workaround: Clear the setting to avoid issues with these IP addresses.

Upgrade might fail if the name of a default port classification has been changed

Description: When you change the original name of a default port classification and then try to upgrade to VMM 1801 – upgrade might fail with the following error message in the VMM setup log.

Violation of PRIMARY KEY constraint 'PK_tbl_NetMan_PortClassification'. Cannot insert duplicate key in object 'dbo.tbl_NetMan_PortClassification'.

Workaround: Change the port classification name back to original name, and then trigger the upgrade. After the upgrade, you can change the default name to a different one.

VMM 2016 release notes

The following sections summarize the release notes for VMM 2016 and includes the known issues, fixes, and the workarounds.

VMM deployment

VMM admin console import might fail

Description: If you import the VMM admin console add-in as a non-administrator, the console will crash. This occurs because the console add-in is stored in location “C:\Program Files\”, and only admins have access to this location. **Workaround:** Store the console add-in at a location that doesn't need admin access, and then import it.

Storage

Promoting a VM to highly available might fail

Description: You create a VM on local storage, start it, and create checkpoints. If you try to migrate and promote the VM as highly available in a cluster, the migration might fail. **Workaround:** Before you run the migration, delete the running checkpoint, and stop the VM.

Migrating a VM from CSV to LUN storage might fail

Description: You create a highly available VM using CSV storage, add a LUN as available storage on the cluster, and migrate the VM from CSV to the LUN. If the VM and LUN storage are on the same node, the migration will

succeed. If they're not, migration will fail. **Workaround:** If the VM isn't located on the cluster node on which the LUN storage is registered, move it there. Then, migrate the VM to the LUN storage.

Capacity of NAS arrays is displayed as 0 GB.

Description: VMM shows **Total Capacity** and **Available Capacity** as 0 GB for existing file shares in the NAS arrays. **Workaround:** None.

Networking

Logical networks managed by SDN network controller can't use dynamic IP addresses

Description: Using dynamic IP addresses for VMs connected to logical networks that are managed by SDN network controller in the VMM fabric isn't supported. **Workaround:** Configure static IP addresses.

SET switch shows as 'Internal' in VMM

Description: If you deploy an SET switch outside the VMM console, and then start managing it in the VMM fabric, The switch type will show as "internal".. This doesn't impact switch functionality. **Workaround:** None

LACP teamed switch doesn't work after upgrade

Description: An LACP team configured in a logical switch doesn't work after upgrading to VMM 2016.

Workaround: Redeploy the switch, or remove and re-add a physical network adapter in the team.

Backend adapter connectivity for SLB MUX doesn't work as expected

Description: Backend adapter connectivity of SLB MUX might not work as expected after VM migration.

Workaround: Use scale in/scale out in the SLB MUX VM as a workaround.

CNG-based CA certificate isn't supported

Description: If you're using certificates from a CA, you can't use CNG certificates for SDN deployment in VMM.

Workaround: Use other certificate formats.

A virtual adapter connected to a network managed by Network Controller must be restarted if you change the IP address

Description: If there's a change in the assigned IP address on any of the virtual network adapters connected to a VM network managed by Network Controller, you need to manually restart the associated adapters. **Workaround:** No workaround.

IPv6 isn't supported for a network infrastructure managed by Network Controller

Description: IPv6 isn't supported by Network Controller in the VMM fabric. **Workaround:** Use IPv4.

Connectivity issues for SLB addresses.

Description: For frontend and backend IP addresses assigned to SLB MUX VMs, you might experience connectivity issues if **Register this connection's address in DNS** is selected. **Workaround:** Clear the setting to avoid issues.

Cluster management

Upgrading the functional level of a cluster doesn't refresh file server information

Description: If you upgrade the functional level for a cluster that includes a file server, the platform information isn't automatically updated in the VMM database. **Workaround:** After upgrading the cluster functional level, refresh the storage provider for the File Server.

Updating a Storage Spaces Direct cluster in VMM will fail

Description: Updating a Storage Spaces Direct cluster (hyper-converged or disaggregated) using VMM isn't supported, and might cause data loss. **Workaround:** Update the clusters outside VMM, using cluster-aware updating (CAU) in Windows.

A cluster rolling upgrade of a Windows Server 2012 R2 host cluster to a Windows Server 2016 Nano Server host cluster will fail

Description: When you try to upgrade the host nodes of a Windows Server 2012 R2 cluster to Windows Server 2016 - Nano Server using cluster rolling upgrade functionality in VMM, the upgrade will fail with error 20406: "VMM could not enumerate instances of class MSFT_StorageNodeToDisk on the server . Failed with error MI RESULT 7 The requested operation is not supported.: **Workaround:** Manually upgrade the Windows Server 2012 R2 host cluster to Nano outside of VMM. Note that rolling upgrade from Windows Server 2012 R2 to Windows Server 2016 Full Server works fine. This issue is specific to Nano.

Adding a cluster in the VMM Admin console might cause an error

Description: When you add a cluster as a resource in the VMM Administrative console, you might receive an error stating "There were no computers discovered based on your inputs". **Workaround:** Select **OK**, and close the error dialog box. Then try to add the cluster again.

A cluster rolling upgrade doesn't do live migration of VMs that aren't highly available

Description: When you run a rolling upgrade of Windows Server 2012 R2 clusters to Windows Server 2016 using VMM, it does not perform live migration of VMs that aren't highly-available. They are moved to a saved state. **Workaround:** Either make all cluster VMs highly available before the upgrade, or do a manual live migration for the specific VMs.

You need manual steps to add a Nano Server-based host located in an untrusted domain

Description: You can't add a Nano Server-based host in an untrusted domain. **Workaround:** Perform these steps on the host and then add it to the VMM fabric as an untrusted host.

1. Enable WINRM over HTTPS:

```
New-Item -Path WSMan:\LocalHost\Listener -Transport HTTPS -Address * -CertificateThumbPrint $cert.Thumbprint -Force
```

2. Create an firewall exception on the host, to allow WINRM over HTTPS:

```
New-NetFirewallRule -DisplayName 'Windows Remote Management (HTTPS-In)' -Name 'Windows Remote Management (HTTPS-In)' -Profile Any -LocalPort 5986 -Protocol TCP
```

You can't add Nano Server-based hosts located in a perimeter network

Description: Trying to add a Nano Server-based host located in a perimeter network using the Add Resource Wizard fails. **Workaround:** Perform these steps on the host, and then add it as an untrusted host to the VMM fabric.

1. Enable WINRM over HTTPS on the host:

```
New-Item -Path WSMan:\LocalHost\Listener -Transport HTTPS -Address * -CertificateThumbPrint $cert.Thumbprint -Force
```

2. Create a firewall exception on the host, to allow WINRM over HTTPS:

```
New-NetFirewallRule -DisplayName 'Windows Remote Management (HTTPS-In)' -Name 'Windows Remote Management (HTTPS-In)' -Profile Any -LocalPort 5986 -Protocol TCP
```

Bare metal deployment of hosts might fail during a highly available upgrade

Description: After a highly available upgrade to VMM 2016, VMM might incorrectly update the Windows Deployment Services (WDS) registry key, HKLM\SYSTEM\CCS\SERVICES\WDSSERVER\PROVIDER\WDSPXE\PROVIDES\VMMOSDPROVIDER, to 'HOST/VIRT-VMM-1', instead of 'SCVMM/VIRT-VMM-1'. This will cause failures in bare metal deployment.

Workaround: Manually change the registry entry for HKLM\SYSTEM\CCS\SERVICES\WDSSERVER\PROVIDER\WDSPXE\PROVIDES\VMMOSDPROVIDER to 'SCVMM/VIRT-VMM-1'.

Host agent status mismatch after upgrade

Description: When VMM updates the host agent, it generates a new certificate for the host. Because of this update, the Network Controller server certificate and the host certificate don't match. **Workaround:** Repairing the host on the **Host Status** page

SAN migration fails for a Nano Server-based host

Description: If you do a SAN migration between two standalone Nano Server-based hosts, an error is issued.

Workaround: Install the latest VMM update rollup (issue fixed in Update Rollup 2)

Storage Spaces Direct

Adding a host with Storage Spaces Direct enabled to the VMM fabric issues a warning

Description: When you add a host to a cluster with Storage Spaces Direct enabled, the warning "Multipath I/O is not enabled for known storage arrays on host <\hostname>" is generated. **Workaround:** Install the latest VMM update rollup (issue fixed in Update Rollup 2).

Deploying a VM on SOFS using fast file copy issues a warning

Description: If you deploy a VM on an SOFS using fast file copy, the action completes successfully with the following warning: "VMM could not transfer the file to using fast file copy. The VMM agent on returned an error

Workaround: None

Cluster validation always runs

Description: When you add a node to a cluster (or create a Storage Spaces Direct hyper-converged cluster), cluster validation is always performed, even when the skip cluster validation option is selected. **Workaround:** Install the latest VMM update rollup. The issue was fixed in Update Rollup 2.

A classification change on a Cluster Shared Volume (CSV) isn't applied

Description: If you change the classification on a Cluster Shared Volume (CSV) in a Storage Spaces Direct hyper-converged cluster, only the classification of the owner node is updated. The other nodes still have the older classification assigned. **Workaround:** Install the latest VMM update rollup. The issue was fixed in update rollup 2.

Creating tiered file share on SOFS doesn't work as expected

Description: When you successfully create a tiered fileshare on SOFS, an error (43020 [SM_RC_DEDUP_NOT_AVAILABLE]) is issued, even if dedup option is not selected. **Workaround:** Ignore the error.

VMM doesn't show correct information for a hyper-converged cluster, or Storage Spaces Direct SOFS

Description: After you add an existing hyper-converged cluster or Storage Spaces Direct SOFS cluster to the VMM fabric, the Storage Provider isn't added, and some properties aren't available. **Workaround:** Install the latest VMM update rollup. The issue was fixed in update rollup 2.

VM management

Shielding a VM causes an error

Description: If you enable shield for an existing VM in the VMM fabric, or if you create a shielded VM from a non-shielded template, the job might fail with the error 1730: "The selected action could not be completed because the virtual machine is not in a state in which the action is valid." The failure happens during the last step of the job, when the VM is shut down after the shielding is completed. The VM is shielded properly, and is usable.

Workaround: Repair the VM with the **Ignore** option.

VMM doesn't show changes to VM security properties

Description: If you change the secure boot properties of a generation 2 VM, or enable/disable vTPM for a shielded VM, outside the VMM console, the change is not immediately shown in VMM. **Workaround:** Manually refresh the VM to show the changes.

Storing a VM in the VMM library fails if you change the default port for BITS (443)

Description: If you change the default BITS port when configuring VMM, an error is issued when you store a VM in the VMM library. Error 2940: "VMM is unable to complete the requested file transfer. The connection to the HTTP Server <name> could not be established." **Workaround:** Manually add the new port number to the Windows Firewall exceptions list of the host:

```
netsh advfirewall firewall add rule name="VMM" dir=in action=allow localport=<port no.> protocol=TCP
```

You can't create VM templates from a Nano Server-based VM.

Description: When you try to create a VM template from a Nano Server-based VM, error 2903 is issued: "VMM could not locate the specified file/folder '' on the '' server. This file/folder might be required as part of another object." **Workaround:** Create a VM template from scratch, using a Nano Server VHD.

Service deployments from service templates might fail on a Nano Server/Core-based guest OS.

Description: When selecting roles and features for a service template, the guest OS profile doesn't differentiate between Core, Nano Server, and Desktop. If you select roles and features (such as Desktop Experience or other GUI-related features) that don't apply to a Core/Nano Server-based guest OS, deployment failure might occur.

Workaround: Don't include these roles and features in the service template.

VMM 2016 doesn't update the VMM guest agents after an upgrade

Description: When you upgrade VMM to 2016 with existing service deployments, and then service those services, VMM 2016 guest agents aren't updated on the VMs that were part of the service deployment. This doesn't affect functionality. **Workaround:** Manually install the VMM 2016 guest agent.

Nano Server-based VM fails to join a domain

Description: During Nano Server VM deployment, if you join the VM to a domain by specifying the domain join information on the **OS Configuration** page of the VM deployment Wizard, VMM deploys the VM but doesn't add it to the specified domain. **Workaround:** After the VM is deployed, manually join the VM to the domain. [Learn more.](#)

Error when starting a VM with Start Ordering

Description: Windows Server 2016 includes the VM Start Ordering feature which defines the order in which dependent VMs are started. This functionality isn't available in VMM, but if you've configured the feature outside the VMM, VMM understands the order in which the VMs will start. However, VMM throws a false positive error (12711): "VMM cannot complete the WMI operation on the server because of an error: [MSCluster_ResourceGroup.Name=] The group or resource is not in the correct state to perform the requested operation." **Workaround:** Ignore the error. The VMs will start in the correct order.

Integration

SQL Server Analysis Services (SSAS) integration doesn't work in VMM and Operations Manager Update Rollup 1.

Description: If you're running Update Rollup 1, you can't configure SSAS for SQL Server. **Workaround:** Download the latest Update Rollups. The issue was fixed in Update Rollup 2.

Next steps

[What's new in Virtual Machine Manager](#)

Manage telemetry settings in VMM

5/30/2019 • 5 minutes to read

This article provides information about how to turn on/off the telemetry settings in System Center - Virtual Machine Manager (VMM).

NOTE

Microsoft does not collect any personal data from the customers. We only listen to events that would help diagnostics in VMM. [Learn more](#)

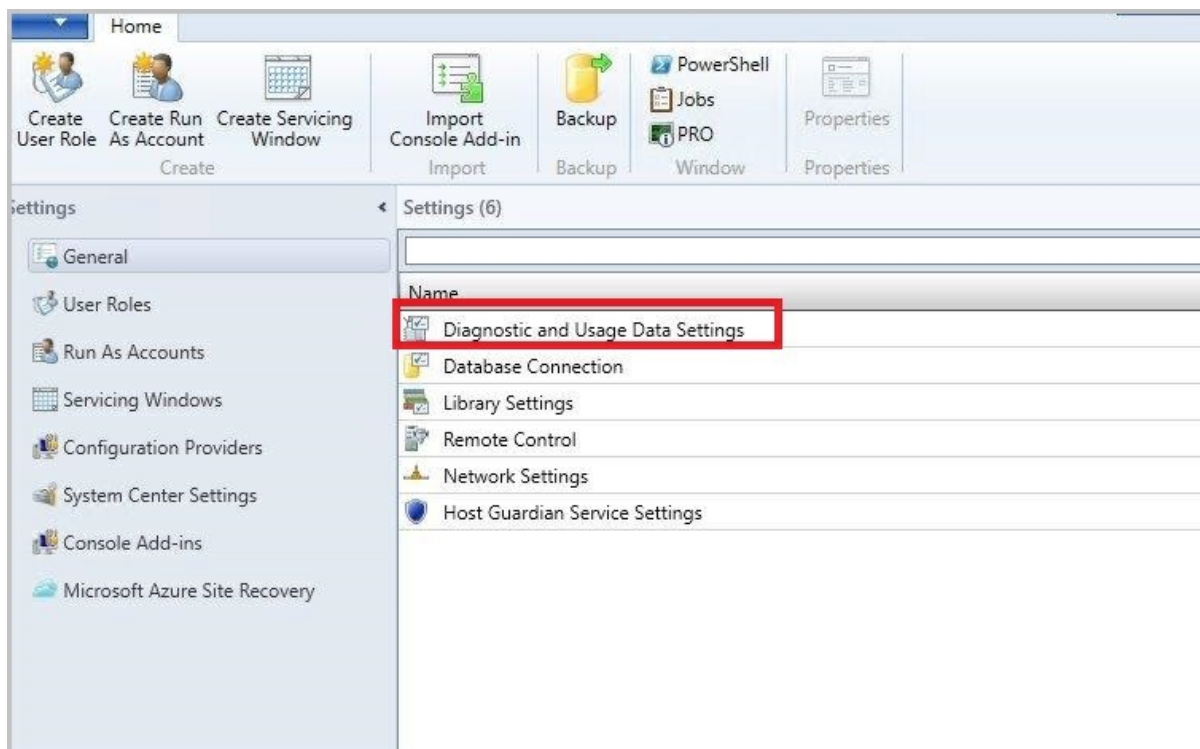
Turn on/off telemetry in VMM

Use the following procedure:

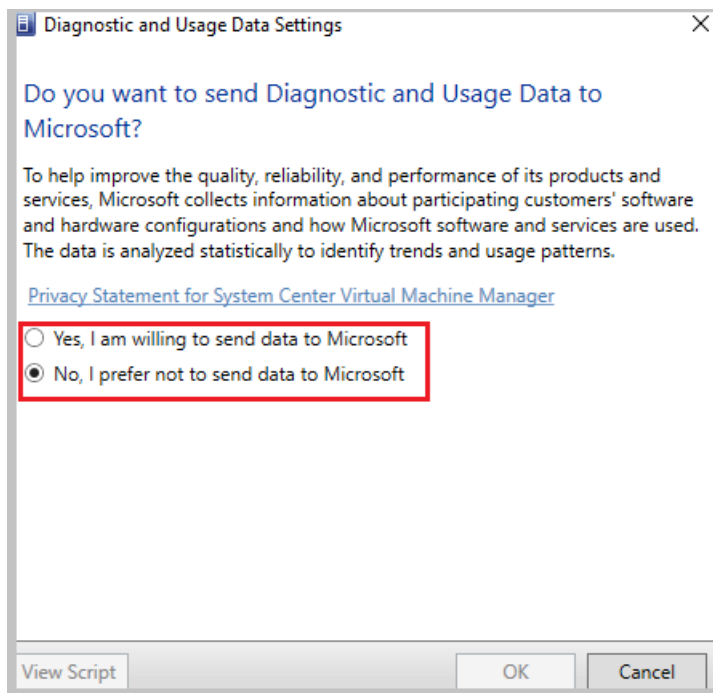
1. On the VMM console, click **Settings**.

General settings window appears with available list of settings.

2. Click **Diagnostics and Usage Data Settings**.



Page with **Do you want to send Diagnostics and Usage Data to Microsoft** message appears with an option to select either Yes or No.



NOTE

We recommend you to go through the privacy statement provided on this page before making a selection.

3. To turn on telemetry, select **Yes I am willing to send data to Microsoft**.
4. To turn off telemetry, select **No I prefer not to send data to Microsoft**.

Telemetry data collected

DATA RELATED TO	DATA COLLECTED
Census Data	Unique ID generated for the VMM deployment VMM build version ID used for correlation with other System Center products

DATA RELATED TO	DATA COLLECTED
Assurance Data	<p>Information about the VMM deployment - whether guardian service deployment or not</p> <p>Count of guarded hosts being managed</p> <p>Count of hosts being managed</p> <p>Count of shielded VM templates</p> <p>Count of VM templates</p> <p>Count of shielded VMs</p> <p>Count of VMs</p> <p>Count of shielded clouds</p> <p>Count of guarded hosts in a cluster</p> <p>Count of guarded hosts in standalone deployments</p> <p>Count of code integrity policy in Global Settings</p> <p>Count of Physical Computer Profile (PCP) with host guarding set</p> <p>If the VMM deployment is a guardian service deployment and OM integrated</p> <p>If the VMM deployment is a guardian service deployment and WSUS integrated</p> <p>Count of hosts with attestation mode as TPM</p> <p>Count of hosts with attestation mode as AD</p> <p>Count of CI policies existing in a library</p> <p>Count of hosts being managed</p> <p>Count of DMZ hosts being managed</p> <p>Count of Hyper-v hosts</p> <p>Count of clusters</p> <p>Count of non-trusted hosts being managed</p> <p>Information if customer is managing the VMWare hosts</p> <p>Minimum, maximum and average host cluster size</p> <p>Minimum, maximum and average host memory size</p> <p>Minimum, maximum and average VMWare cluster size</p>

DATA RELATED TO	DATA COLLECTED
Job related	<p>Indication if jobs are user initiated or background jobs</p> <p>Count of the number of times a job description type has failed</p> <p>Number of times job has run</p> <p>Maximum, minimum and average number of jobs run per day</p>
Dynamic Optimization Settings	<p>Count of hosts considered for dynamic optimization</p> <p>Count of hosts considered for power optimization</p> <p>Minimum, maximum and average frequency to run dynamic optimization</p>
Network	<p>Count of Edge NAT instances</p> <p>Count of Edge forwarding instances</p> <p>Count of Edge S2S instances</p> <p>Count of VPN connections</p> <p>Count of tenants per multi -tenant gateway</p> <p>Count of Edge load balancers</p> <p>Count of Edge load balancer VIP templates</p> <p>Count of NC managed VPN connections</p> <p>Count of NC managed IPsec VPN connections</p> <p>Count of NC managed GRE VPN connections</p> <p>Count of NC managed L3 VPN connections</p>
IP Pool	<p>Count of IPV4 and IPV6 address pools</p> <p>Maximum size of IPV4 and IPV6 pools</p> <p>Total reserved IPV4 and IPV6 addresses</p> <p>Average pool usage size</p> <p>Count of NC managed IPV4 and IPV6 addresses</p>
Library	<p>Count of library shared being managed</p> <p>Count of VHDs and VHDXs registered in the library</p> <p>Count of service templates registered in the library</p> <p>Count of VM templates registered in the library</p> <p>Count of host groups belonging to the library</p> <p>Count of equivalent resources registered in the library</p>

DATA RELATED TO	DATA COLLECTED
Logical Network	<p>Count of Logical networks</p> <p>Count of logical network definitions</p> <p>Count of one connected network managed by network controller</p> <p>Count of VLAN based independent networks</p> <p>Count of private VLAN networks</p>
Policy Distribution	<p>Minimum and maximum policy RTT</p> <p>Policy errors</p>
Service Objects	<p>Count of service instances</p> <p>Count of deployed service instances</p> <p>Count of failed service instances</p> <p>Count of servicing failed service instances</p> <p>Count of service instances deployed to non Hyper-V host groups</p> <p>Count of VM roles</p> <p>Count of applications</p> <p>Count of script applications</p> <p>Maximum, minimum and average tiers per service</p> <p>Count of services using application hosts</p> <p>Count of services using SQL application</p> <p>Count of services using Web deploy application</p> <p>Maximum, minimum and average number of applications per customer, computer tier</p> <p>Count of service tiers with load balancers</p> <p>Count of services deployed on VMware</p>

DATA RELATED TO	DATA COLLECTED
VM	<p>Count of HA VM instances</p> <p>Count of Non HA VM instances</p> <p>Count of HA VM instances on clustered storage</p> <p>Count of HA VM instances on SMB share</p> <p>Count of DiffDisk VM instances</p> <p>Count of Hyper-V VM instances</p> <p>Count of VMWare VM instances</p> <p>Count of OS type, OS name, OS edition, OS version on VMs</p> <p>List of possible VM states</p> <p>Count of VMs per state</p>
VM OS	<p>OS version of VMs</p> <p>OS name of the VMs</p> <p>OS edition of VMs</p> <p>OS Type of VMs</p> <p>Count of VMs for various OS</p> <p>List of possible VM states</p> <p>Count of VMs for each state</p>
VMM Settings	<p>Indication of highly available installation of VMM</p> <p>Indication if VMM is upgraded</p> <p>VMM SKU type</p> <p>Indication of VMM running as domain account</p> <p>Indication if VMM DB is on remote server</p> <p>Indication if VMM is running on VM</p> <p>Indication if VMM is running on default path</p> <p>Count of VM networks</p> <p>Count of VM subnets</p> <p>Count of NC managed VM networks and VM subnets</p>

DATA RELATED TO	DATA COLLECTED
Host OS	<p>OS version on the host</p> <p>OS name of the host</p> <p>SKU type of host</p> <p>Indication of Nano type host</p> <p>Number of hosts with specific OS</p>
Storage Inventory	<p>Count of all storage providers</p> <p>Count of SMIS CIMXML, SMIS WMI, Native Windows, SMP WMI storage providers</p> <p>Count of providers managing arrays</p> <p>Count of providers managing file shares</p> <p>Count of providers managing both arrays and file shares</p> <p>Count of providers managing FC fabric</p> <p>Count of storage classification defined by user</p> <p>Count of discs where custom classification is used to override inherited disk classification</p> <p>Count of disks where custom classification is used to override pool derived classification</p> <p>Count of fabric classifications</p> <p>Count of primary VMs for which HVR is enabled out of band</p> <p>Count of VMs for which HVR is enabled in band</p> <p>Count of LUN based protected replication groups created Out of Band</p> <p>Count of volume-based protected replication groups created Out of Band</p> <p>Count of volume-based protected replication groups created through Azure Site Recovery</p> <p>Indication if VMM is configured to connected to Azure Site Recovery</p> <p>Count of QoS policies created in VMM</p> <p>Count of VMs with QoS policies applied</p>

DATA RELATED TO	DATA COLLECTED
Storage Hyper Converged cluster	<p>Count of Hyper converged clusters</p> <p>Count of pools</p> <p>Count of virtual discs</p> <p>Count of volumes</p> <p>Count of local discs attached</p> <p>Average local disc size attached</p>
SAN Arrays	<p>Count of pools in storage array</p> <p>Count of managed pools in storage array</p> <p>Count of managed LUNs in all managed pools per array</p> <p>List of array capabilities</p> <p>Indication if array also has associated file share</p>
Storage Hosts	<p>Count of storage hosts or clusters</p> <p>Name of storage aggregate collection</p> <p>Maximum, minimum, average, median and mode values of storage aggregate of collector</p>
File server storage	<p>List of total file shares on file server</p> <p>List of managed shares on file server</p> <p>Indication if there is calabria file share capability</p> <p>List of associated storage array capabilities</p> <p>List of SMB file share dialect supported by fileserver</p>
Telemetry Opt out	<p>Indication if Telemetry is opted out</p>

Scenario: Deploy a private VMM cloud

12/7/2018 • 2 minutes to read

This article provides an overview of System Center - Virtual Machine Manager (VMM) private clouds.

A private cloud is a cloud that is provisioned and managed on-premises by an organization. It's deployed using an organization's own hardware to leverage the advantages of a private cloud model.

You can use VMM to create and deploy private cloud components, and to manage access to the private cloud and the underlying physical resources. VMM provides the following benefits:

- **Self-service**—Self-service admins can delegate management and usage of the private cloud, while having no knowledge of underlying physical resources. They don't have to ask the private cloud provider for administrative changes, except to request increased capacity and quotas as required.
- **Opacity**—Self-service admins don't need any knowledge of underlying physical resources.
- **Resource pooling**—Administrators can collect and present an aggregate set of resources, such as storage and networking resources. Resource usage is limited by the capacity of the private cloud and by user role quotas.
- **Elasticity**—Administrators can add resources to a private cloud to increase the capacity.
- **Optimization**—Usage of the underlying resources is continually optimized without affecting the overall private cloud user experience.

You can create a private clouds from either:

- VMM host groups that contain resources from virtualization hosts.
- A VMware resource pool

You deploy a private cloud by configuring fabric resources, settings up library paths for private cloud users, and setting cloud capacity.

Next steps

- [Create a private cloud in VMM](#)

Scenario: Create a private cloud

5/6/2019 • 10 minutes to read

This article provides instructions for creating a private cloud in System Center - Virtual Machine Manager (VMM).

You can create a private cloud from a host group, or from a VMware resource pool. Host groups can contain a single host type or a mix of Hyper-V and VMware ESX hosts.

NOTE

You must be a VMM administrator, or a member of the delegated Administrators user group, with a scope that includes the host groups you'll use for the cloud.

Before you start

- You need to have the VMM fabric in place. [Learn more](#).
- You should have one or more [Hyper-V](#) or [VMware](#) virtualization hosts in the fabric. If you're creating a cloud from a VMware resource pool, a vCenter Server and the VMware ESX host or host cluster that contains the VMware resource pool must be available in the VMM fabric.
 - If you want to provide [self-service users](#) the ability to store virtual machines to the VMM library, then create a library share, or create a folder in a library share that will serve as the storage location. [Learn more](#).

NOTE

Self-service users must have the **Store and re-deploy** permission to store their virtual machines.

- You can add independently created library shares as read-only library shares in the VMM when you run the Create Cloud Wizard. For example, outside the VMM, you cannot create the **\\VMMServer01\Finance\StoredVMs** path, and then add the **\\VMMServer01\Finance** library share to the VMM library.
- The self-service user role data path is specified when you create a self-service user role or modify the properties of a self-service user role.
- If you want to assign read-only shares to the private cloud, where administrators can store read-only resources such as .iso files that they want to make available to self-service users, make sure that one or more library shares exists that you can assign as the read-only library shares.

NOTE

Self-service users must have the **Author** permission to access the resources.

- The library shares that you designate as read-only resource locations for the private cloud must be unique when compared to the library share or shares that are used for stored virtual machines and for the user role data path that is specified for a self-service user role. For example, if the user role data path for a self-service user role is **\\VMMServer01\Finance**, you cannot create a stored virtual machine path of **\\VMMServer01\Finance\StoredVMs**. However, if the user role data path is **\\VMMServer01\Finance\FinanceUserRoleData**, you could specify

\\VMMServer01\Finance\StoredVMs as the stored virtual machine path, as the full path is unique. You could also create entirely separate library shares.

Create a private cloud from a host group

1. Click **VMs and Services > Create > Create Cloud**, to open the Create Cloud Wizard.
2. In **General**, specify a **Name** and optional description for the cloud.
3. Specify whether the cloud will support shielded VMs.
4. In **Resources > Host groups**, select the groups you want to add to the cloud. Then click **Next**.
5. In **Logical Networks**, select each logical network that you want to make available to the private cloud, and then click **Next**.

Note: Only logical networks that are associated with the physical network adapters on hosts in the selected host groups appear in the list.

6. In **Load Balancers**, select each load balancer that you want to make available to this private cloud, and then click **Next**.

Note: Only load balancers that are associated with the selected host groups appear in the list.

7. In **VIP Templates**, select each VIP template that you want to make available to the private cloud, and then click **Next**.

8. In **Port Classifications**, select each port classification that you want to make available to the cloud, and then click **Next**.

9. In **Storage**, if you have storage managed by VMM, select each storage classification that you want to make available to the private cloud, and then click **Next**.

Note: Only storage classifications for storage pools that are assigned to the selected host groups appear in the list.

10. In **Library > Stored VM path**, browse and select the library share you want to use for the self-service users to store VMs. Click **OK**.

11. In **Read-only library shares > Add**, select one or more library shares where administrators can provide read-only resources to cloud users. Click **OK** and then click **Next**.

12. In **Capacity**, set capacity limits for the private cloud, and then click **Next**. You can either accept the default values, or clear the **Use Maximum** check boxes and set quotas for the following resources:

QUOTA TYPE	DESCRIPTION
Virtual CPUs	Sets a limit on processing capacity within the private cloud that is equivalent to the capacity that can be provided by a specified number of CPUs. Applied against running virtual machines. Setting a CPU quota does not guarantee contiguous capacity; it only guarantees total CPU capacity available among hosts in the private cloud.

QUOTA TYPE	DESCRIPTION
Memory	Sets a quota on memory (in gigabytes) that is available for virtual machines that are deployed on the private cloud. Applied against running virtual machines only. Setting a memory quota does not guarantee contiguous capacity. For example, the private cloud might have 2 GB of memory available on one host and 2 GB of memory on the other.
Storage	Sets a quota on storage capacity (in gigabytes) that is available to virtual machines that are deployed on the private cloud. For dynamic virtual hard disks, quota calculations are based on maximum size.
Custom quota (points)	Sets a quota on virtual machines that are deployed on the private cloud based on the total quota points that are assigned to the virtual machines through their virtual machine templates. Quota points are an arbitrary value that can be assigned to a virtual machine template based on the anticipated size of the virtual machines. Custom quotas are provided for backward compatibility with self-service user roles that were created in VMM.
Virtual machines	Limits the total number of virtual machines that can be deployed on the private cloud.

13. In **Capability Profiles**, select each virtual machine capability profile that you want to add, and then click **Next**. Select the capability profiles that match the type of hypervisor platforms that are running in the selected host groups. The built-in capability profiles represent the minimum and maximum values that can be configured for a virtual machine for each supported hypervisor platform.
14. In **Replication Groups**, select the replication groups for the private cloud, and click **Next**.
15. In **Summary** page, confirm the settings, and then click **Finish**.

View status in **Jobs** and ensure the job is complete.

To verify that the private cloud was created, check **VMs and Services > Clouds**. You can also verify in **Library > Cloud Libraries**, to view the read-only library shares.

Assign Storage QoS Policies while creating a cloud

1. Follow the steps until 14 in the [above procedure](#).
2. In **Storage QoS Policies**, select the policies that you want to assign to this cloud.
3. Proceed with rest of the steps and complete the wizard.

Create a private cloud from a VMware resource pool

1. Click **VMs and Services > Create > Create Cloud**, to open the Create Cloud Wizard.
2. In **General**, specify a **Name** and optional description for the cloud.
3. In **Resources**, select **VMware resource pools > Next**.
4. In **Logical Networks**, select each logical network that you want to make available to the private cloud, and then click **Next**.
5. In **Load Balancers**, select each load balancer that you want to make available to the private cloud, and then

click **Next**. Only load balancers that are associated with the selected host groups appear in the list.

6. In **VIP Templates**, select each VIP template that you want to make available to the private cloud, and then click **Next**.
7. In **Storage**, click **Next**. VMM doesn't manage or assign storage classifications assigned to ESX hosts.
8. In **Library > Stored VM path**, browse and select the library share you want to use for self-service users to store VMs, click **OK**.
9. In **Read-only library shares > Add**, select one or more library shares where administrators can provide read-only resources to cloud users. Click **OK** and then click **Next**.
10. In **Capacity**, set capacity limits for the private cloud, and then click **Next**. You can either accept the default values, or clear the **Use Maximum** check boxes and set quotas for the following resources:

QUOTA TYPE	DESCRIPTION
Virtual CPUs	Sets a limit on processing capacity within the private cloud that is equivalent to the capacity that can be provided by a specified number of CPUs. Applied against running virtual machines. Setting a CPU quota does not guarantee contiguous capacity; it only guarantees total CPU capacity available among hosts in the private cloud.
Memory	Sets a quota on memory (in gigabytes) that is available for virtual machines that are deployed on the private cloud. Applied against running virtual machines only. Setting a memory quota does not guarantee contiguous capacity. For example, the private cloud might have 2 GB of memory available on one host and 2 GB of memory on the other.
Storage	Sets a quota on storage capacity (in gigabytes) that is available to virtual machines that are deployed on the private cloud. For dynamic virtual hard disks, quota calculations are based on maximum size.
Custom quota (points)	Sets a quota on virtual machines that are deployed on the private cloud based on total quota points that are assigned to the virtual machines through their virtual machine templates. Quota points are an arbitrary value that can be assigned to a virtual machine template based on the anticipated size of the virtual machines. Custom quotas are provided for backward compatibility with self-service user roles that were created in VMM.
Virtual machines	Limits the total number of virtual machines that can be deployed on the private cloud.

11. In **Capability Profiles**, select **ESX Server**, and then click **Next**. The built-in capability profiles represent the minimum and maximum values that can be configured for a virtual machine for each supported hypervisor platform.
12. In **Summary** page, confirm the settings, and then click **Finish**.

View status in **Jobs** and ensure the job is complete.

To verify that the private cloud was created, check **VMs and Services > Clouds**. You can also verify in **Library > Cloud Libraries**, to view the read-only library shares.

Assign storage QoS policies to a private cloud

System Center - Virtual Machine Manager (SCVMM) 1801 and later, support storage QoS policies on a private cloud.

The VMM fabric admin can now offer the storage QoS policies in the cloud. Tenant admins and self-service users can consume these while deploying the VMs and services. This will enable the cloud providers to guarantee and/or limit the amount of storage performance as per the subscription opted by the tenants.

The fabric admin can now offer storage QoS policies, while authoring the VMM private clouds. After which, the authorized users (admin and self-service users) with access to this cloud can consume the available QoS policies for provisioning VMs and Services on the cloud.

NOTE

A cloud with at least one QoS policy is offered, will not allow the deployment of disks without a policy. This helps in preventing the cloud consumers to pick infinite resources with a null policy. Also, the change in IOPS by self-service users will not be allowed after this release, to avoid the same scenario of selecting unauthorized performance.

Storage QoS Policies option in the **Create cloud** wizard helps the fabric admin to select the list of policies, which should be made available for the cloud consumers.

follow these steps:

1. Follow the [create a private cloud from a host group](#) procedure until step 14.
2. In **Storage QoS Policies**, select the policies that you want to assign to this cloud.
3. proceed to step 15, complete the remaining steps.
4. Review the summary and click **Finish**.

NOTE

This list contains only those policies that are available in scope of all the clusters selected in the Resources section of the cloud wizard. This helps the self-service user to choose between the available plans, even after the VM is placed .

On Upgrade

1. After the upgrade, the existing clouds will not have any QoS policy in their offering. Admin needs to update the cloud with the policy offerings.
2. The existing VMs on the cloud, which have disks with a QoS policy already assigned, will go to inconsistent state. Their policy stays intact, but the VMM UI displays it as blank. Admins can either remove those policies, or offer these in the affected clouds.
3. Proceed with rest of the step and complete the wizard.

Assign private clouds to user roles

After you create a private cloud, you can assign the private cloud to one or more user roles.

1. In **VMs and Services**, click the private cloud that you want to assign.
2. Right-click the selected cloud, click **Assign Cloud**.
3. Select an existing user role, or click **Create a user role and assign this cloud** to create a new one.

Manage VMM cloud settings

12/7/2018 • 2 minutes to read

This article describes how to manage System Center - Virtual Machine Manager (VMM) cloud settings.

Manage cloud capacity

You place virtual machines in a VMM cloud if they fit within your capacity settings for the cloud.

- By default VMM assumed that all resources allocated to a replica VM are in use, and uses this assumption when figuring out whether placing a VM in a cloud or host group will fit within the cloud or host group limits.
- If you want to override this default behavior, you can configure a registry key that allows you to overcommit cloud and host groups in the VMM fabric. In other words, you can placing a VM in a cloud or host group even if that placement will bring a cloud or group above its capacity limits. This is useful if you know that you don't need all your replica VMs to be running simultaneously.

Example with the default settings

You have a cloud with a maximum setting of 32 GB of memory. That cloud have two VMs with 4 GB of memory each. If you not want to place a third VM with 26 GB of memory the action will fail with "Not enough memory available"

Example with the registry setting

You have a cloud with a maximum setting of 32 GB of memory. That cloud have two VMs with 4 GB of memory each. To place a third VM with 26 GB of memory you do the following:

1. Update the capacity in the cloud properties to accommodate the increased memory size. Say to 64 GB.
2. Navigate to registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings\Placement\IgnoreMemoryForStoppedReplicaVM**.
3. Set the DWORD value to 1. If the value **IgnoreMemoryForStoppedReplicaVM**.doesn't exist, create it.

Scenario - Deploy guarded hosts and shielded virtual machines in VMM

7/9/2019 • 2 minutes to read

This article provides an overview of deploying Hyper-V guarded hosts and shielded virtual machines in a System Center - Virtual Machine Manager (VMM) compute fabric.

Guarded fabrics provide additional protections for VMs to prevent tampering and theft by malicious administrators and malware. As a cloud service provider or private cloud administrator, you can deploy a guarded fabric that typically consists of a server running the host guardian service (HGS), one or more guarded Hyper-V host servers, and one or more shielded VMs running on those hosts. [Learn more about guarded fabrics.](#)

Why do I need to protect VMs?

Virtual machines contain sensitive data and configuration that the VM owner may not want a fabric administrator to see. However, since all the data for VMs are stored in files, the data can easily be copied off and inspected by malware or a malicious administrator. Shielded VMs in Windows Server help prevent such attacks by rigorously attesting to the health of a Hyper-V host before booting up a VM, ensuring the VM can only be started in datacenters authorized by the VM owner, and enabling the guest OS to encrypt its own data through the use of a new, virtual TPM. The VM owner can select from the following two types of protection when creating a security-sensitive VM:

- **Encryption Supported:** Ideal for enterprise private cloud scenarios where encryption of data at rest and in-flight is necessary, but the fabric administrators are still trusted. The VM console and other management conveniences remain available to fabric administrators.
- **Shielded:** The most secure deployment option, shielding prevents fabric administrators from connecting to the VM console or modifying security aspects of the VM configuration. VM owners can only access the VM through remote management tools they choose to enable. This is recommended for tenants running sensitive workloads on public or shared infrastructure.

Managing a guarded fabric with VMM

The core guarded fabric infrastructure (consisting of one or more guarded Hyper-V hosts, the Host Guardian Service, and the artifacts needed to create shielded VMs) is included with Windows Server 2016 and above, and must be configured according to the [guarded fabric documentation](#). Once set up, you can optionally use System Center - Virtual Machine Manager to simplify management of the guarded fabric.

VMM can be used to:

- **Provision and manage guarded hosts in the VMM fabric:** You can add and manage guarded hosts to the VMM fabric. A guarded host is a Hyper-V server that:
 - Meets the [guarded host prerequisites](#).
 - Is authorized by the Host Guardian Service for the fabric to run shielded VMs. The HGS admin determines the requirements for hosts to successfully attest and become "guarded".
 - Is marked as guarded in VMM by configuring it to use the same HGS URLs as those specified in the global VMM settings.
- **Configure a shielded virtual hard disk and optionally a VM template:** Signed template disks (VHDX) used to deploy new shielded VMs can be stored in the VMM library for easy deployment. You can then use this VHDX in a VM template.

- **Provision and manage shielded VMs:** VMM supports the full lifecycle of shielded VMs. This includes:
 - Creating new shielded VMs from a signed template disk (VHDX), and optionally using a VM template.
 - Converting existing VMs to shielded VMs.

Next steps

- [Provision guarded hosts in the VMM fabric](#)

Provision guarded hosts in VMM

12/7/2018 • 6 minutes to read

This article describes how to deploy guarded Hyper-V hosts in a System Center - Virtual Machine Manager (VMM) compute fabric. [Learn more](#) about guarded fabric..

There are a couple of ways to set up guarded Hyper-V hosts in a VMM fabric.

- **Configure an existing host to be a guarded host:** You can configure an existing host to run shielded VMs.
- **Add or provision a new guarded host:** This host could be:
 - An existing Windows Server computer (with or without the Hyper-V role)
 - A bare-metal computer

You set up guarded hosts in the VMM fabric as follows:

1. **Configure global HGS settings:** VMM connects all guarded hosts to the same HGS server so that you can successfully migrate shielded VMs between the hosts. You specify global HGS settings that apply to all guarded hosts, although you can specify host-specific settings that override the global settings. Settings include:
 - **Attestation URL:** The URL that the host uses to connect to the HGS attestation service. This service authorizes a host to run shielded VMs.
 - **Key protection server URL:** The URL that the host uses to retrieve the key needed to decrypt VMs. The host must pass attestation in order to retrieve keys.
 - **Code integrity policies:** A code integrity policy restricts the software that can run on a guarded host. When HGS is configured to use TPM attestation, guarded hosts must be configured to use a code integrity policy authorized by the HGS server. You can specify the location of code integrity policies in VMM and deploy them to your hosts. This is optional and not required to manage a guarded fabric.
 - **VM shielding helper VHD:** A specially-prepared virtual hard disk that is used to convert existing VMs to shielded VMs. You must configure this setting if you wish to shield existing VMs.
2. **Configure the cloud:** If the guarded host will be included in a VMM cloud, you need to enable the cloud to support shielded VMs.

Before you start

Make sure you have deployed and configured the Host Guardian Service before proceeding. [Learn more](#) about configuring HGS in the Windows Server documentation.

Additionally, ensure any hosts which will become guarded hosts meet the guarded host prerequisites:

- **Operating system:** Host servers must run Windows Server Datacenter. It is recommended to use Server Core for guarded hosts.
- **Role and features:** Host servers should be running the Hyper-V role and the Host Guardian Hyper-V Support feature. Host Guardian Hyper-V Support lets the host communicate with HGS to attest to its health and request keys for shielded VMs. If your host is running Nano Server, it should have the Compute, SCVMM-Package, SCVMM-Compute, SecureStartup, and ShieldedVM packages installed.
- **TPM-attestation:** If your HGS is configured to use TPM attestation, the host servers must:
 - Use UEFI 2.3.1c and a TPM 2.0 module
 - Boot in UEFI mode (not BIOS or "legacy" mode)
 - Enable Secure Boot

- **HGS registration:** Hyper-V hosts must be registered with HGS. How they're registered depending on whether HGS is using AD or TPM attestation. [Learn more](#)
- **Live migration:** If you want to live migrate shielded VMs, you need to deploy two or more guarded hosts.
- **Domain:** Guarded hosts and the VMM server must be in the same domain, or in domains with a two-way trust.

Configure global HGS settings

Before you can add guarded hosts to your VMM compute fabric, you must configure VMM with information about the Host Guardian Service for the fabric. The same HGS will be used for all guarded hosts managed by VMM.

1. Obtain the attestation and key protection URLs for your fabric from your HGS administrator.
2. In the VMM console, click **Settings > Host Guardian Service Settings**.
3. Enter the attestation and key protection URLs in the respective fields. You do not need to configure the code integrity policies and VM shielding helper VHD sections at this time.

Host Guardian Service Settings

Specify Host Guardian Service and related settings

When you enable the Host Guardian Service on a host, the following URLs will be configured on the host. Virtual Machine Manager (VMM) will also use the Key Protection Server URL as the destination for tenant keys.

Attestation Server URL:

 Example: http://contoso.com/Attestation

Key Protection Server URL:

 Example: http://contoso.com/KeyProtection

Code Integrity policies

These policies restrict the software that can run on the host. After you click Finish, when you configure properties for a host, you can select the policies to apply to that host, if those policies are in a path that is accessible by the host computer account.

Policies:

Name	File Path
HW1-SC1-Enforced	\\vmm.contoso.com\\CodeIntegrityPolicies\\HS1-SC1-Enfo...

Buttons: Add, Remove

Shielding Helper VHD

Select a VHD that VMM will use for OS volume encryption while shielding existing virtual machines.

Buttons: Browse..., Clear

Buttons: View Script, Finish, Cancel

4. Click **Finish** to save the configuration.

Add or provision a new guarded host

1. Add the host:
 - If you want to add an existing server running Windows Server as a guarded Hyper-V host, [add it to the fabric](#).
 - If you want to provision a Hyper-V host from a bare-metal computer, [follow these prerequisites and instructions](#). Note that you can deploy the host as guarded when you provision it (Add Resource Wizard

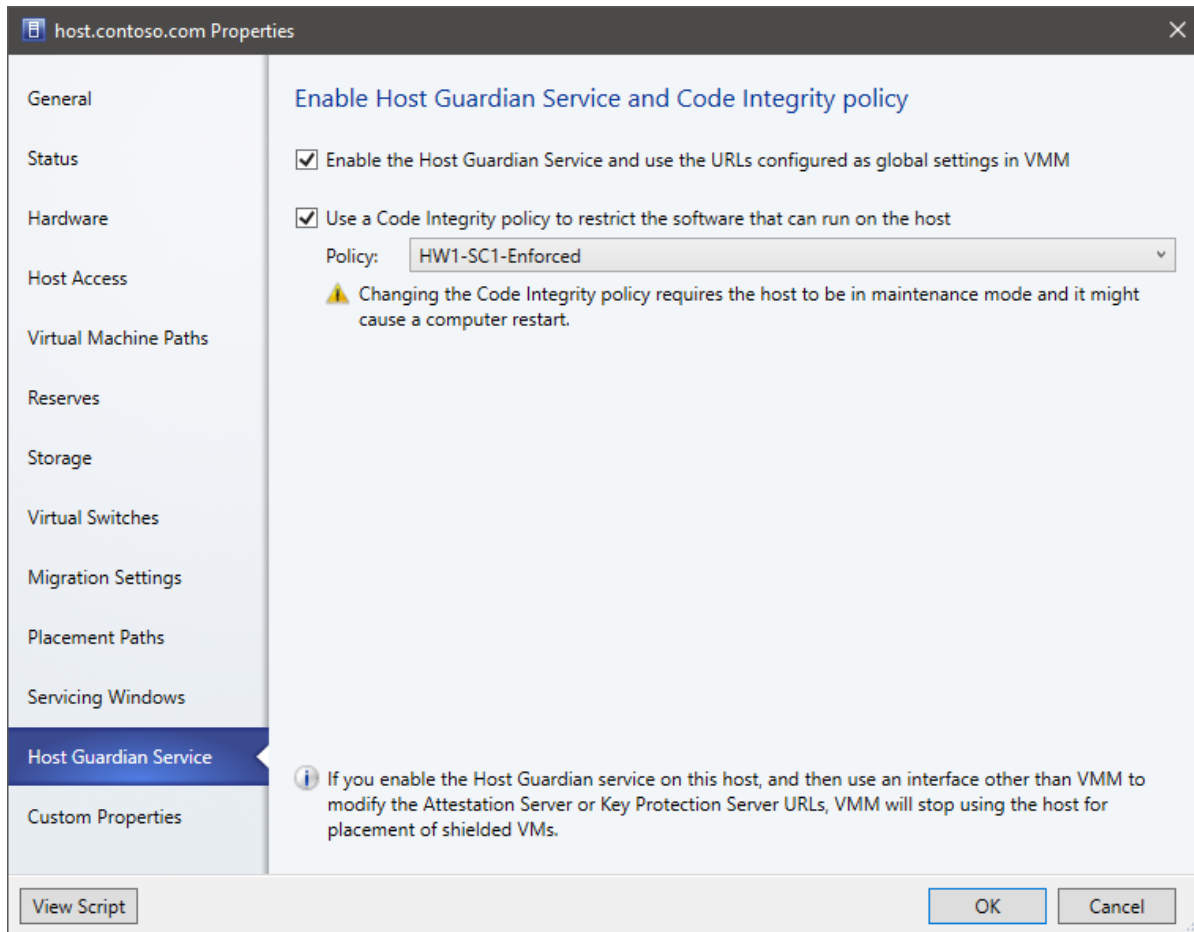
> **OS Settings** > **Configure as guarded host.**

2. Continue on to the next section to configure the host as a guarded host.

Configure an existing host to be a guarded host

To configure an existing Hyper-V host managed by VMM to be a guarded host, complete the following steps:

1. Place the host in [maintenance mode](#).
2. In **All Hosts**, right-click the host > **Properties** > **Host Guardian Service**.



3. Select to enable the Host Guardian Hyper-V Support feature and configure the host. Note that:
 - The global attestation and key protection server URLs will be set on the host.
 - If you modify these URLs outside the VMM console, you need to update them in VMM too. If you don't, VMM will not place shielded VMs on the host until the URLs match again. You can also uncheck and re-check the "Enable" box to reconfigure the host with the URLs configured in VMM.
4. If you're using VMM to manage code integrity policies, you can enable the second checkbox and select the appropriate policy for the system.
5. Click **OK** to update the host's configuration.
6. Take the host out of maintenance mode.

VMM checks that the host passes attestation when you add it, and every time that the host status is refreshed. VMM only deploys and migrates shielded VMs on hosts that have passed attestation. You can check the attestation status of a host in **Properties** > **Status** > **HGS Client Overall**.

Enable guarded hosts on a VMM cloud

Enable a cloud to support guarded hosts:

1. In the VMM console click **VMs and Services** > **Clouds**. Right-click the cloud name > **Properties**.
2. In **General** > **Shielded VM support**, select **Supported on this private cloud**.

Manage and deploy code integrity policies with VMM

In guarded fabrics configured to use TPM attestation, each host must be configured with a code integrity policy that is trusted by the Host Guardian Service. To ease the management of code integrity policies, you can optionally use VMM to deploy new or updated policies to your guarded hosts.

To deploy a code integrity policy to a guarded host managed by VMM, complete the following steps:

1. [Create a code integrity policy](#) for each reference host in your environment. You will need a different CI policy for each unique hardware and software configuration of your guarded hosts.
2. Store the CI policies in a secure file share. The computer accounts for each guarded host require **read access** to the share. Only trusted administrators should have write access.
3. In the VMM console, click **Settings** > **Host Guardian Service Settings**.
4. Under the Code Integrity Policies section, click **Add** and specify a friendly name and the path to a CI policy. Repeat this step for each unique CI policy. Be sure to name your policies in a manner that will help you identify which policy should be applied to which hosts.

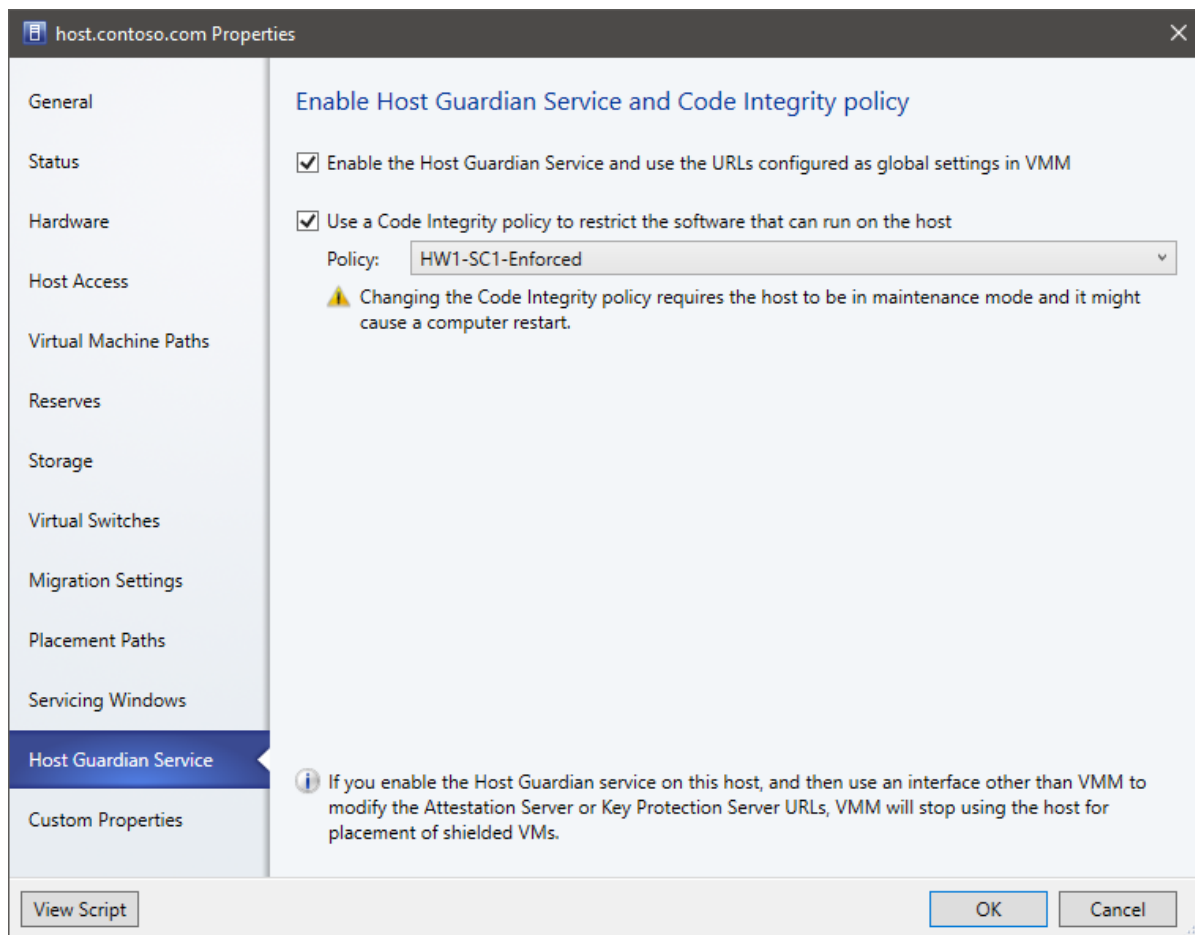
The screenshot shows the 'Host Guardian Service Settings' dialog box. It has a title bar with a shield icon and the text 'Host Guardian Service Settings'. The main content area is titled 'Specify Host Guardian Service and related settings'. Below this, there is explanatory text: 'When you enable the Host Guardian Service on a host, the following URLs will be configured on the host. Virtual Machine Manager (VMM) will also use the Key Protection Server URL as the destination for tenant keys.' There are two text input fields: 'Attestation Server URL:' with the value 'http://hgs.secure.contoso.com/Attestation' and an example 'http://contoso.com/Attestation'; and 'Key Protection Server URL:' with the value 'http://hgs.secure.contoso.com/KeyProtection' and an example 'http://contoso.com/KeyProtection'. Below these is a section for 'Code Integrity policies' with explanatory text: 'These policies restrict the software that can run on the host. After you click Finish, when you configure properties for a host, you can select the policies to apply to that host, if those policies are in a path that is accessible by the host computer account.' There is a 'Policies:' section with a table. The table has two columns: 'Name' and 'File Path'. It contains one row: 'HW1-SC1-Enforced' and '\\vmm.contoso.com\CodeIntegrityPolicies\HS1-SC1-Enfo...'. To the right of the table are 'Add' and 'Remove' buttons. Below the table is a section for 'Shielding Helper VHD' with explanatory text: 'Select a VHD that VMM will use for OS volume encryption while shielding existing virtual machines.' There is a text input field with the value '\\VMM.contoso.com\MSSCVMLibrary\VHDs\ShieldingHelperVHD.vhdx' and 'Browse...' and 'Clear' buttons. At the bottom of the dialog are 'View Script', 'Finish', and 'Cancel' buttons.

Name	File Path
HW1-SC1-Enforced	\\vmm.contoso.com\CodeIntegrityPolicies\HS1-SC1-Enfo...

5. Click **Finish** to save the configuration.

Now, for each guarded host, complete the following steps to apply a code integrity policy:

1. Place the host in [maintenance mode](#).
2. In **All Hosts**, right-click the host > **Properties** > **Host Guardian Service**.



3. Select to enable the option to configure the host with a code integrity policy, then select the appropriate policy for the system.
4. Click **OK** to apply the configuration change. The host may restart to apply the new policy.
5. Take the host out of maintenance mode.

WARNING

Be sure to select the correct code integrity policy for the host. If an incompatible policy is applied to the host, some applications, drivers, or operating system components may no longer work.

If you update the code integrity policy in the file share and wish to also update the guarded hosts, you can do so by completing the following steps:

1. Place the host in [maintenance mode](#).
2. In **All Hosts**, right-click the host > **Apply Latest Code Integrity Policy**.
3. Take the host out of maintenance mode.

Next steps

- [Set up a shielded template disk, utility disk, and VM template](#)

Configure HGS fallback URLs in VMM

5/30/2019 • 2 minutes to read

This article describes how to define the fallback Host Guardian Service (HGS) URLs in System Center - Virtual Machine Manager (VMM) global settings. For information about guarded fabrics, see [this article](#).

Being at the heart of providing attestation and key protection services to run shielded VMs on Hyper-V hosts, the host guardian service (HGS) should operate even in situations of disaster.

With fallback HGS configuration feature in VMM, a guarded host can be configured with a primary and secondary pair of HGS URLs (an attestation and key protection URI). This capability will enable scenarios such as guarded fabric deployments spanning two data centers for disaster recovery purposes, HGS running as shielded VMs etc.

The primary HGS URLs will always be used in favor of the secondary. If the primary HGS fails to respond after the appropriate timeout and retry count, the operation will be reattempted against the secondary. Subsequent operations will always favor the primary; the secondary will only be used when the primary fails.

Before you start

Ensure you have deployed and configured the Host Guardian Service before proceeding. Learn more about [configuring HGS](#).

Configure fallback HGS

Use the following steps:

1. Navigate to **VMM Settings > General Settings > Host Guardian Service settings**.
2. Open the Host Guardian Service settings. You should see a section for Fallback Configurations.
3. Define the **primary and fallback HGS URLs** and click **Finish**.

Host Guardian Service Settings

Specify Host Guardian Service and related settings

When you enable the Host Guardian Hyper-V Support on a host, the following URLs will be configured on the host. Virtual Machine Manager (VMM) will also use the Key Protection Server URL as the destination for tenant keys.

Primary Configurations

Attestation Server URL:

Example: http://contoso.com/Attestation

Key Protection Server URL:

Example: http://contoso.com/KeyProtection

Fallback Configurations

When either the attestation client or the key protection client using the primary host guardian service (HGS) URLs indicates a failure due to connection level incident (connection timeout, reset or other non-recoverable error), the HGS client on the host restarts the requested operation using the specified Fallback HGS URLs.

Not all the Hyper-V hosts support the HGS Fallback configuration. To know about the supported hosts, see [HGS documentation](#).

Attestation Server URL:

Example: http://contoso.com/Attestation

Key Protection Server URL:

Example: http://contoso.com/KeyProtection

Code Integrity policies

These policies restrict the software that can run on the host. After you click Finish, when you configure properties for a host, you can select the policies to apply to that host, if those policies are in a path that is accessible by the host computer account.

Policies:

Name	File Path

Add

Remove

Shielding Helper VHD

Select a VHD that VMM will use for OS volume encryption while shielding existing virtual machines.

Browse...

Clear

View Script

Finish

Cancel

4. Enable the fallback URLs on the host by navigating to **Host Properties > Host Guardian Service** > select Enable host Guardian Hyper-V support and use the URLs as configured as global settings in VMM and click OK.

NOTE

After this step, VMM service configures the supported hosts with primary and fallback HGS URLs. Only hosts on and above Windows Server 1709 support fallback HGS URLs.

PowerShell command updates

1. The following two parameters are added to the existing **Set-SCVMHost** PowerShell command:

- **AttestationFallbackServerUrl**
- **KeyProtectionFallbackServerUrl**

Here is the sample syntax.

```

Set-SCVMHost [-VMHost] <Host> [-ApplyLatestCodeIntegrityPolicy] [-AttestationServerUrl <String>]
[-AttestationFallbackServerUrl <String>]
[-AvailableForPlacement <Boolean>] [-BMCAAddress <String>]
[-BMCCustomConfigurationProvider <ConfigurationProvider>] [-BMCPort <UInt32>]
[-BMCPProtocol <OutOfBandManagementType>] [-BMCRunAsAccount <RunAsAccount>] [-BaseDiskPaths <String>]
[-BypassMaintenanceModeCheck] [-CPUPercentageReserve <UInt16>] [-CodeIntegrityPolicy
<CodeIntegrityPolicy>]
[-Custom1 <String>] [-Custom10 <String>] [-Custom2 <String>] [-Custom3 <String>] [-Custom4 <String>]
[-Custom5 <String>] [-Custom6 <String>] [-Custom7 <String>] [-Custom8 <String>] [-Custom9 <String>]
[-Description <String>] [-DiskSpaceReserveMB <UInt64>] [-EnableLiveMigration <Boolean>]
[-FibreChannelWorldWideNodeName <String>] [-FibreChannelWorldWidePortNameMaximum <String>]
[-FibreChannelWorldWidePortNameMinimum <String>] [-IsDedicatedToNetworkVirtualizationGateway <Boolean>]
[-JobGroup <Guid>] [-JobVariable <String>] [-KeyProtectionServerUrl <String>] [-
KeyProtectionFallbackServerUrl <String>] [-LiveMigrationMaximum <UInt32>]
[-LiveStorageMigrationMaximum <UInt32>] [-MaintenanceHost <Boolean>] [-ManagementAdapterMACAddress
<String>]
[-MaxDiskIOReservation <UInt64>] [-MemoryReserveMB <UInt64>]
[-MigrationAuthProtocol <MigrationAuthProtocolType>]
[-MigrationPerformanceOption <MigrationPerformanceOptionType>] [-MigrationSubnet <String[]>]
[-NetworkPercentageReserve <UInt16>] [-NumaSpanningEnabled <Boolean>] [-OverrideHostGroupReserves
<Boolean>]
[-PROTipID <Guid>] [-RemoteConnectCertificatePath <String>] [-RemoteConnectEnabled <Boolean>]
[-RemoteConnectPort <UInt32>] [-RemoveRemoteConnectCertificate] [-RunAsynchronously] [-SMBiosGuid
<Guid>]
[-SecureRemoteConnectEnabled <Boolean>] [-UseAnyMigrationSubnet <Boolean >]
[-VMHostManagementCredential <VMMCredential>] [-VMPaths <String>] [<CommonParameters>]

```

- The following parameter is added to **Get-SCGuardianConfiguration** to let the user specify from which HGS the metadata be fetched.

[-Guardian {Primary | Fallback}]

Syntax

```

Get-SCGuardianConfiguration [-Guardian {Primary | Fallback}] [-OnBehalfOfUser <String>] [-
OnBehalfOfUserRole <UserRole>] [-VMMServer <ServerConnection>] [<CommonParameters>]

```

Next steps

- [Deploy the Host Guardian Service \(HGS\)](#)
- [Manage HGS](#)
- [Set up a fallback HGS for a branch office](#)

Set up a disks and a VM template to deploy shielded VMs

12/7/2018 • 4 minutes to read

You deploy shielded virtual machines in the System Center - Virtual Machine Manager (VMM) compute fabric using a signed virtual machine hard disk (VHDX), and optionally with a VM template. This article describes how to add signed template disks to VMM, configure a shielding utility disk, deploy new shielded VMs and convert existing VMs to shielded VMs in VMM.

Before you start

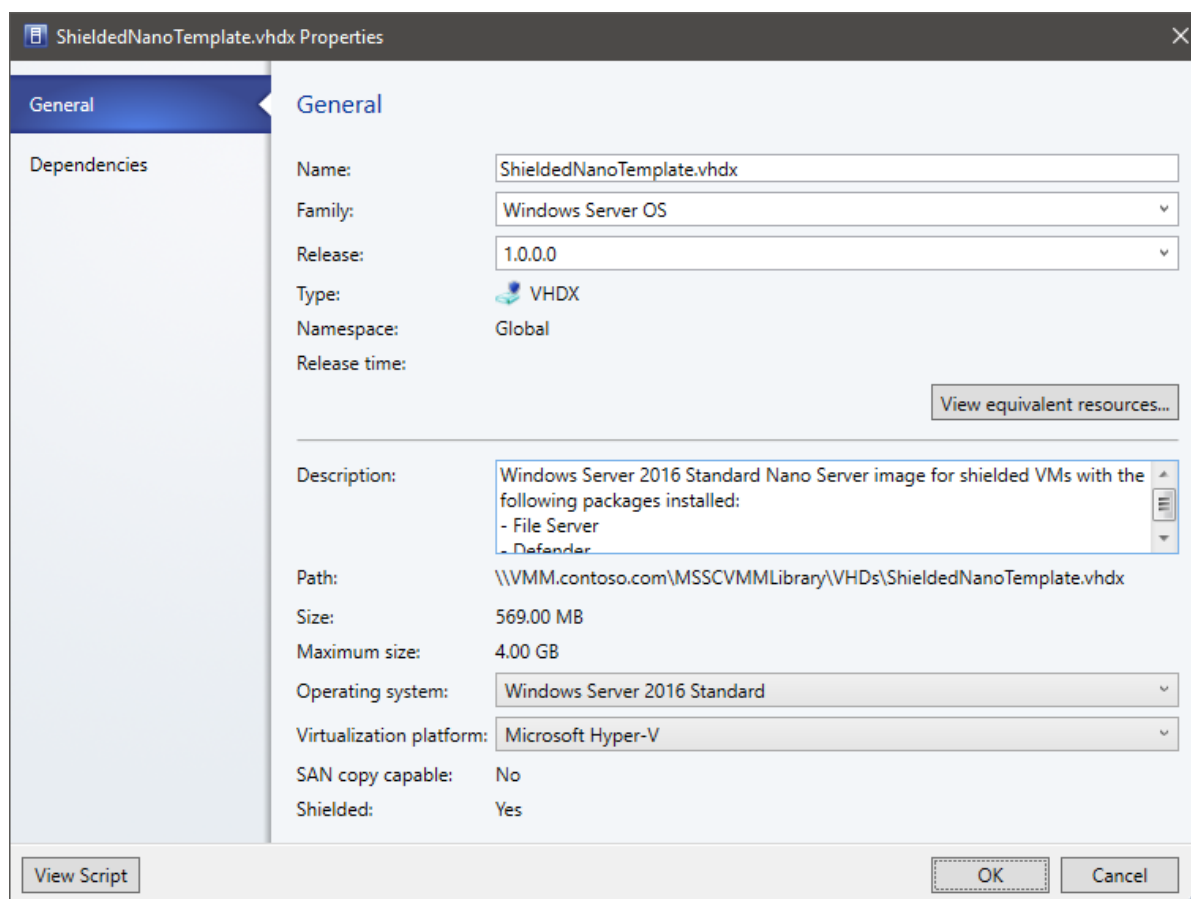
- The signed template disk used to create the shielded VM template must have the family and version marked.
- The VMM library to which you add the signed template disk must be accessible to clouds from which shielded VMs will be provisioned.
- The library shared should be added to clouds from which shielded VMs will be provisioned (not in read-only mode).

Adding signed template disks for shielded VMs to the VMM library

Shielded VMs can be deployed in two ways: by deploying directly from a signed template disk or by converting an existing VM to a shielded VM. Signed template disks assure tenants that the disk contents have not been modified and enable tenants to securely transfer deployment secrets like administrator passwords and certificates to the VM in an encrypted manner. For this reason, it is preferred to deploy shielded VMs from signed template disks.

To prepare and add a signed template disk to the VMM library, complete the following steps:

1. [Prepare a signed template disk](#) on a machine running Windows Server 2016 with Desktop Experience, or Windows 10 with the [Remote Server Administration Tools](#) installed.
2. [Copy the template disk](#) to a library share (\\<vmmserver>\MSSCVMMLibrary\VHDs by default), and refresh the library server.
3. To provide VMM with information about the operating system on the template disk, in **Library**, right-click the disk > **Properties**.
4. In **Operating system**, select the operating system installed on the disk. This indicates to VMM that the VHDX isn't blank. The shield icon next to the disk name denotes it as a signed template disk for shielded VMs. Supply information about the **Family** and **Release** of the disk as well to make the resources available in the tenant Azure Pack self-service portal (optional).



5. Click **OK** to save the properties of the signed template disk.

Create a shielded VM template

You can optionally create a shielded VM template using a signed template disk. VM templates define virtual machine resources such as CPU count, RAM, and networking for an OS disk.

Templates for shielded VMs vary slightly from a regular VM template. Some settings are fixed – for example the VM must be a Generation 2 VM with Secure Boot enabled. Create the VM template as follows:

1. Click **Library > Create VM Template**. In **Select Source**, click Use an existing VM template or a virtual hard disk stored in the library > **Browse**.
2. Select the signed template disk, specify a template name and optional description, and click **OK**.
3. In **Configure Hardware**, specify the hardware properties for VMs you create from the template. Make sure there's at least one NIC configured and available. Tenants connect to shielded VMs over Remote Desktop Connection, Windows Remote Management, or other remote management tools that require networking.
4. If you want to use static IP addressing in the tenant pool, you need to let your tenants know. Tenants need to provide an answer file with values that specializes a shielded VM for them. There are special, well-known placeholder values required to support static IP pools.
5. In **Configure Operating System**, specify the OS version, computer name, product key and time zone. The tenant provides secure information such as the administrator password in a shielding data file (.PDK) that they'll provide when provisioning a new VM. If you specify a product key make sure it's valid for the operating system on the template disk. If it isn't, the VM will not provision successfully. After the VM template is created, make sure that it's available to the Tenant Administrator user role. Tenants can then use it to provision new VMs.

Configure the shielding helper VHD

Existing Windows VMs can also be converted to shielded VMs with the use of a shielding helper VHD. The helper VHD is a special disk prepared with tools to encrypt another VM's operating system drive. VMM must be

configured with a helper VHD before you can shield existing VMs.

1. [Prepare a helper VHD](#) on a computer running Windows Server 2016 or Windows 10 with the [Remote Server Administration Tools](#) installed.
2. [Copy the helper VHD](#) to a library share, and refresh the library server.
3. In the VMM console, click **Settings > Host Guardian Service Settings**.
4. In the Shielding Helper VHD section, click **Browse** and select the helper VHD from the list of files in the library shares.
5. Click **Finish** to save the configuration.

With the shielding helper VHD configured, you can proceed to [shield an existing VM](#).

Next steps

Review [Provision shielded VMs](#) to understand how to deploy shielded virtual machines in a VMM compute fabric.

Provision shielded virtual machines in the VMM fabric

5/30/2019 • 4 minutes to read

This article describes how to deploy shielded virtual machines in the System Center - Virtual Machine Manager (VMM) compute fabric.

You can deploy shielded VMs in VMM in a couple of ways:

- Convert an existing VM into a shielded VM.
- Create a new shielded VM using a signed virtual machine hard disk (VHDX), and optionally a VM template.

Before you start

[Watch](#) a video that provides a quick, two-minute overview of provisioning shielded VMs in VMM. Then, make sure you've done the following:

1. **Prepare an HGS server:** You should have an HGS server deployed. [Learn more](#).
2. **Set up VMM:** You need to configure global HGS settings in VMM, and set up at least one guarded host. If guarded hosts belong to a cloud, the cloud should be enabled to support shielded VMs. [Learn more](#).
3. **Prepare a shielded VHDX and VM template:** You deploy shielded VMs from a shielded virtual hard disk (VHDX), optionally using a VM template. [Learn more](#) about preparing these.

NOTE

You cannot use a service template to create a shielded VM. Use a script instead.

4. **Prepare shielding data files:** To use the signed template disks in the VMM library, tenants must prepare one or more shielding data files. This file contains all the secrets that a tenant needs to deploy a VM, including the unattend file used to specialize the VM, certificates, administrator account passwords. The file also specifies which guarded fabric a tenant trusts to host their VM and information about the signed template disks. The file is encrypted and can only be read by a host in a guarded fabric trusted by the tenant. [Learn more](#).
5. **Set up host group:** For easy management, we recommend that guarded hosts be placed in a dedicated VMM host group.
6. **Verify existing VM requirements:** If you want to convert an existing VM to shielded, note the following:
 - The VM must be generation 2 and have the Microsoft Windows Secure Boot template enabled
 - The operating system on the disk must be one of:
 - Windows Server 2016, Windows Server 2012 R2, Windows Server 2012
 - Windows 10, Windows 8.1, Windows 8
 - The OS disk for the VM must use GUID Partition Table. This is required for generation 2 VMs to support UEFI.
7. **Set up helper VHD:** The hosting service provider will need to create a VM that acts as a helper VHD for converting existing machines. [Learn more](#).

Adding shielding data files to VMM


Before you can convert an existing VM to a shielded VM or provision a new shielded VM from a template, the VM owner must generate a shielding data file and add it to VMM.

If you do not already have a shielding data file imported, complete the following steps:

1. [Create a shielding data file](#) if you don't already have one. Make sure the shielding data file authorizes the hosting fabric VMM manages to run your shielded VMs.
2. In the VMM console, click **Library > Import Shielding Data > Browse** and select your shielding data file.
3. Specify a friendly name for the shielding data file in **Name** and optionally add a description. It is recommended that you indicate whether the shielding data file is intended for use with existing or new VMs in its name to make it easier to find again.
4. Click **Import** to save the shielding data in VMM.

To manage your imported shielding data files, go to **Library > VM Shielding Data** (under "Profiles").

Provision a new shielded VM

1. Make sure you have all the prerequisites in place before you start.
2. In **VMs and Services**, click **Create Virtual Machine** to open the Create Virtual Machine Wizard.
3. In **Select Source**, click **Use an existing virtual machine, VM template, or virtual hard disk > Browse**.
4. Select a shielded VM template or signed template disk. Both are identified by the shield icon .
5. In **Select Shielding Data File**, click **Browse** and select a shielding data file. Only shielding data files that can be used to create a new shielded VM will be shown. Click **OK > Next** to continue.
6. Follow [these instructions](#) to complete the wizard, and to deploy the VM on a host/cloud.

When you complete the wizard, VMM creates a new shielded VM from the disk or template:

1. The template disk (VHDX) file is copied from the VMM library
2. VM provisioning decrypts the data in the shielding data file, completes any substitution strings in the unattend.xml file, and copies additional files from the shielding data file to the operating system drive (for example, the RDP certificate).
3. The VM restarts, is customized, and re-encrypted with BitLocker. The BitLocker full volume encryption key is stored in the virtual TPM of the new VM.
4. VM customization is complete when the shutdown command in the unattend.xml file runs, the VM remains switched off. If customization gets stuck, check the unattend.xml file by running it on an unshielded VM, or using an encryption-supported shielding data file that allows console access.
5. After VMM detects that specialization has finished, it will update its status to indicate the VM is created and, if selected, start up the VM.

Shield an existing VM

You can enable shielding for a VM currently running on a host in the VMM fabric that isn't guarded.

1. Ensure you have all the prerequisites in place before you start.
2. Take the VM offline.
3. We recommend that you enable BitLocker on all disks attached to the VM before moving it to the guarded host.
4. Select the VM > **Properties > Shield**, and select a shielding data file.
5. Shut down the VM, export from non-guarded host, and import it to a guarded host. Only a guarded host can access the VM data.

Next steps

Review [Manage virtual machine settings](#) to learn how to configure performance and availability settings for VMs.

Provision a shielded Linux virtual machine in the VMM fabric

5/30/2019 • 3 minutes to read

This article describes how to deploy Linux shielded virtual machines (VMs) in System Center 1801 - Virtual Machine Manager (VMM).

Procedure to shield a Linux VM

Windows Server 2016 introduced the concept of a shielded VM for Windows OS-based virtual machines. Shielded VMs provide protection against malicious administrator actions when the VM's data is at rest or when untrusted software is running on Hyper-V hosts. [Learn more](#).

With Windows Server version 1709, Hyper-V introduced support for provisioning Linux shielded VMs. This support is available in VMM 1801.

Shield a Linux VM

1. Create a signed template disk.
2. Create a Linux shielded VM template in VMM.
3. Generate a shielding data file (PDK).
4. Create a Linux shielded VM by using the VM template and the PDK.

NOTE

If you use Wireless Application Protocol (WAP), you can provision Linux shielded VMs in the same way you provision Windows shielded VMs.

Prepare a template disk

1. Follow [these steps](#) to create the template disk.
2. In the **Preparing a Linux Image** section of the directions, before you install lsvmttools, [install the VMM specialization agent](#).

Sign the template disk

1. Generate a certificate. You can use a self-signed certificate for testing.

Use the following sample cmdlet:

```
$cert = New-SelfSignedCertificate -DnsName '<<signing.contoso.com>>'
```

2. Sign the disk by using a Windows Server 1709 machine. Use the following sample cmdlet:

```
Protect-TemplateDisk -Path "<<Path to the VHDX>>" -TemplateName "<<Template Name>>" -Version <<x.x.x.x>>  
-Certificate $cert -ProtectedTemplateTargetDiskType PreprocessedLinux
```


3. Copy the template disk and the signed image to the VMM library.

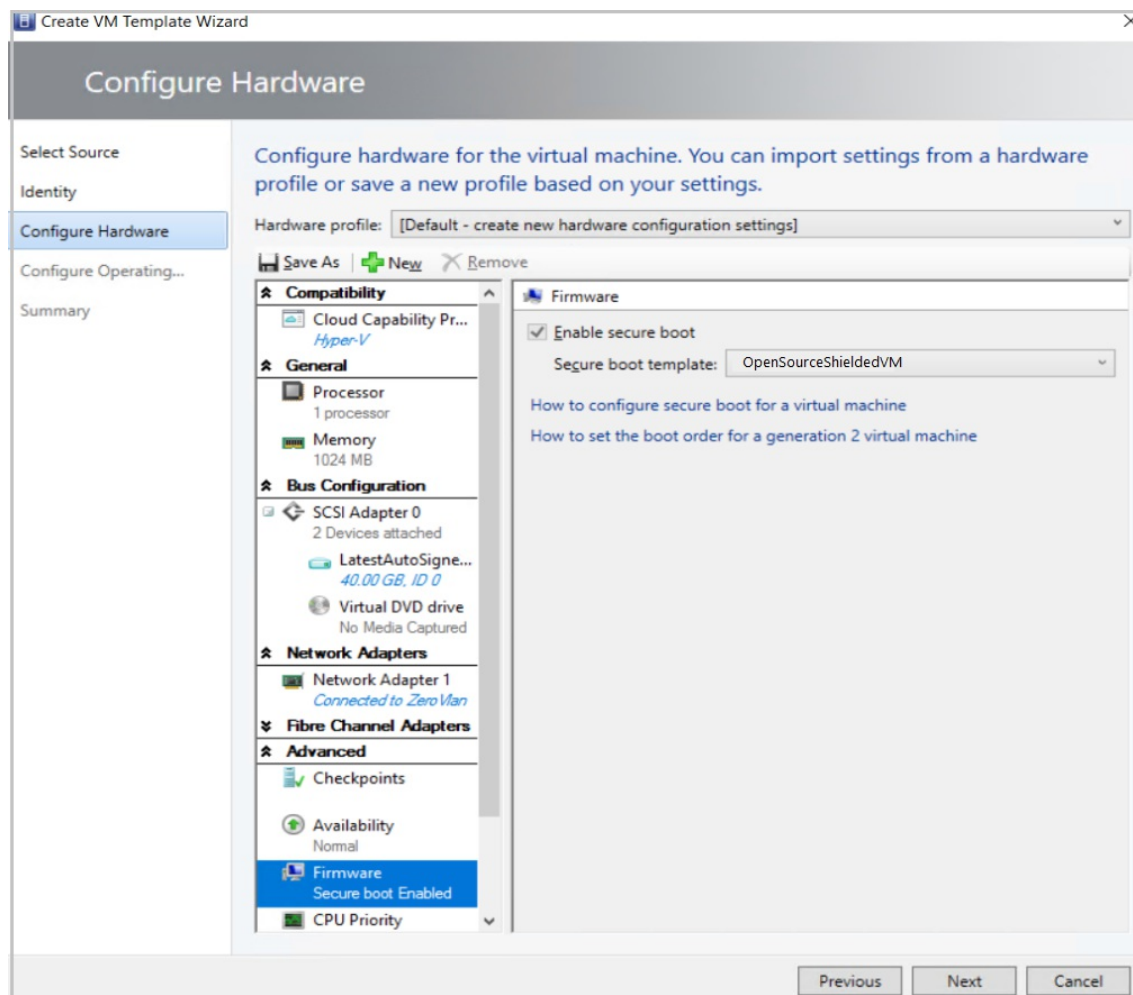
Create a Linux shielded VM template in VMM

1. In the VMM console library, select **Create VM Template**.
2. In **Select Source**, select **Use an existing VM template**. Browse to select the signed template disk that you added to the VMM library. Then select **Next**.
3. In **Configure Hardware**:
 - Under **Firmware**, select **Enable secure boot**. From the **Secure boot template** drop-down menu, select **OpenSourceShieldedVM**.

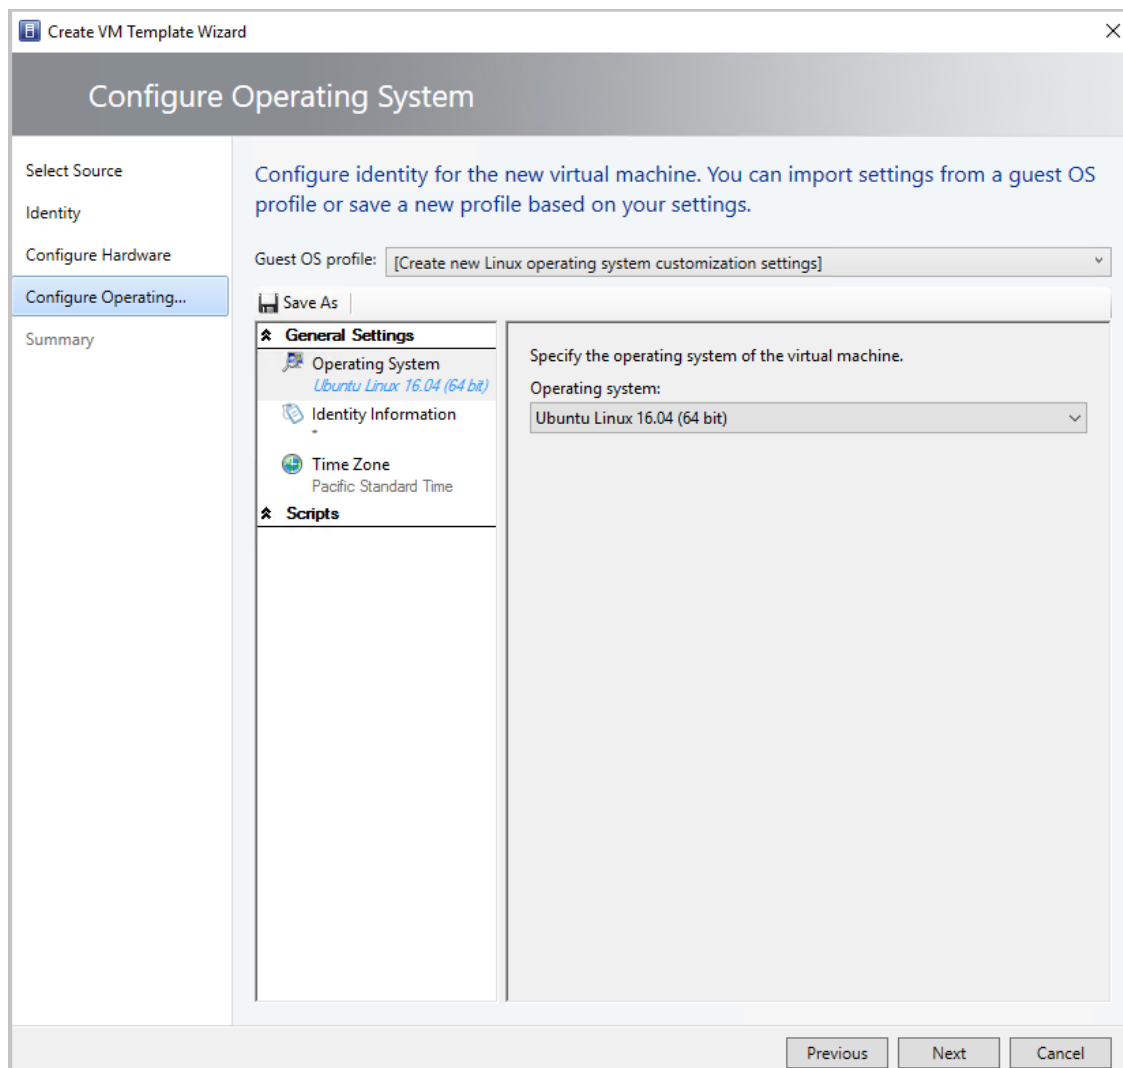
NOTE

This boot template is a new addition to RS3 hosts. If no RS3 hosts are in VMM, this option won't show up on the **Secure boot template** menu.

- Select the required configuration for other hardware properties, such as processors, memory, and the VM network.



4. In **Configure Operating System**:
 - Select the Guest OS profile as **Create new Linux operating system customization settings**.
 - Select the OS on the template disk that you created earlier (**Ubuntu Linux**).



5. Select **Next**.
6. In **Summary**, review the details and select **Create to finish generation of Linux shielded VM template in VMM**.

Generate the shielding data file

Before you generate the shielding data file (PDK):

1. [Get the guardian metadata from the Host Guardian Service \(HGS\)](#).
2. [Extract the volume signature catalog \(VSC\) file](#).

To generate the PDK, run the following sample script on a server that's running Windows Server version 1709:

```
# Create a VolumeSignatureCatalog file for the template disk to ensure that no one tampers with the template
disk at the deployment time
# Create an owner certificate
$Owner = New-HgsGuardian -Name '<<Owner>>' -GenerateCertificates

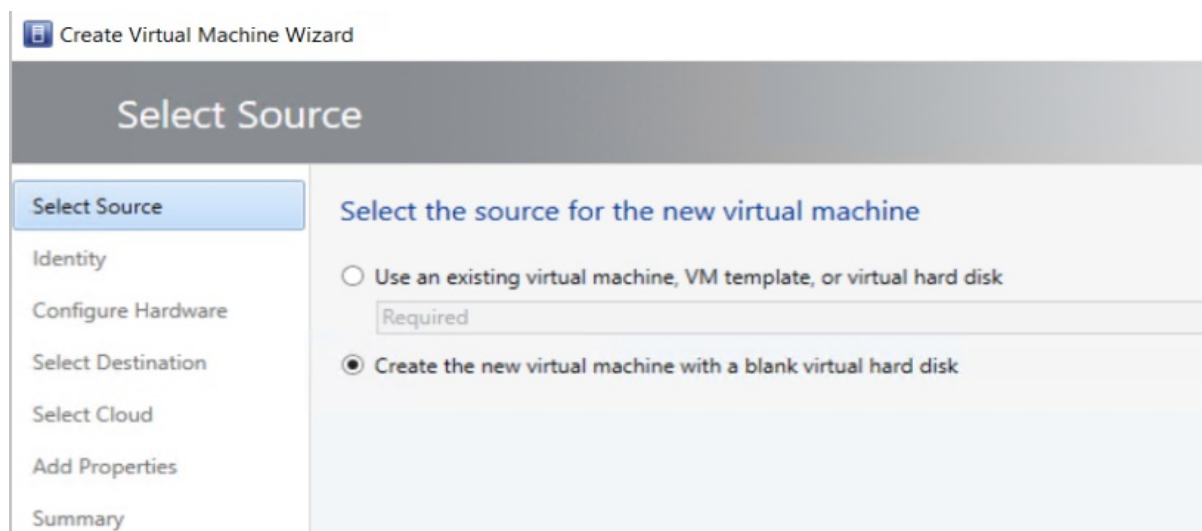
# Import the HGS guardian
$Guardian = Import-HgsGuardian -Path <<Import the xml from pre-step 1>> -Name '<<Name of the guardian>>' -
AllowUntrustedRoot

# Create the PDK file on a server running Windows Server version 1709

New-ShieldingDataFile -ShieldingDataFilePath '<<Shielding Data file path>>' -Owner $Owner -Guardian $guardian -
VolumeIDQualifier (New-VolumeIDQualifier -VolumeSignatureCatalogFilePath '<<Path to the .vsc file generated in
pre-step 2>>' -VersionRule Equals) -AnswerFile '<<Path to LinuxOsConfiguration.xml>>' -policy Shielded
```

Create a Linux shielded VM by using the VM template and the PDK

1. In the VMM console, select **Create Virtual Machine**.
2. Select **Use an existing virtual machine, VM template, or virtual hard disk**.
3. Select **Linux shielded VM template > Next**.



4. Name the VM and select **Next**.
5. In **Configure Hardware**, make sure the details match your template settings. Then select **Next**.
6. In **Configure Operating System** settings, ensure the details conform to the settings you made when you created the template. Then select **Next**.
7. Select the shielding data file (PDK) that you created.
8. Select the destination host group and then select **Next**.
9. Select the host by the rating that the VMM placement engine gave. Then select **Next**.
10. In **Configure Settings**, review the virtual machine settings and select **Next**.
11. Review the actions in **Add properties** and select **Next**.
12. To create the Linux shielded VM, select **Create**.

While provisioning the VM, the VMM specialization agent reads the Linux configuration file PDK and customizes the VM.

Next steps

- Get an overview of [Guarded fabric and shielded VMs](#).
- Find out more about [Linux shielded VM tools](#).

Deploy and manage a Software Defined Network (SDN) infrastructure in the VMM fabric

12/7/2018 • 2 minutes to read

System Center - Virtual Machine Manager (VMM) can be used to deploy and manage a Software Defined Network (SDN) infrastructure.

Software Defined Network overview

SDN virtualizes your network to abstract physical hardware network elements such as switches and routers. Using SDN you can dynamically manage your datacenter networking to meet workload and app requirements. Network policies can be implemented consistently, at scale, even as you deploy new workloads, or move workloads across virtual or physical networks.

If you deploy SDN in the VMM fabric you can:

- Provision and manage virtual networks at scale.
- Deploy and manage the SDN infrastructure, including network controllers, software load balancers, and gateways.
- Define and control virtual network policies centrally and link them to your applications or workloads. When your workload is deployed or moved, the network configuration adjusts itself automatically. This is important because it removes the need for manual reconfiguration of network hardware, thereby reducing operational complexity while saving your valuable resources for higher-impact work.
- Control traffic flow between virtual networks, including the ability to define guaranteed bandwidth for your critical applications and workloads.

SDN combines a number of technologies, among them:

- **Network Controller:** The network controller allows you to automate configuration of your network infrastructure, instead of manually configuring network devices and services.
- **RAS Gateway for SDN:** RAS Gateway is a software-based, multitenant, BGP capable router in Windows Server 2016 that is designed for CSPs and Enterprises that host multiple tenant virtual networks using HNV.
- **Software Load Balancing (SLB) for SDN:** (SDN) in Windows Server 2016 can use Software Load Balancing (SLB) to evenly distribute tenant and tenant customer network traffic among virtual network resources. The Windows Server SLB enables multiple servers to host the same workload, providing high availability and scalability.

[Read more about](#) technologies in the SDN stack.

Next steps

- [Deploy SDN components using PowerShell](#)
- Alternatively, deploy SDN components manually in the VMM console:
 - [Set up a network controller](#)
 - [Set up a software load balancer](#)
 - [Set up a RAS gateway](#)

Set up an SDN network controller in the VMM fabric

5/30/2019 • 22 minutes to read

This article describes how to set up a Software Defined Networking (SDN) network controller in the System Center - Virtual Machine Manager (VMM) fabric.

The SDN network controller is a scalable and highly available server role that enables you to automate network infrastructure configuration instead of performing manual network device configuration. [Learn more.](#)

For a great introduction, [watch a video](#) (~ five minutes) that provides an overview of network controller deployment.

Prerequisites

- Plan for a Software Defined Network (SDN). [Learn more.](#)
- Plan for an SDN Network Controller Installation and deployment. [Learn more.](#)

Before you start

To set up SDN in the VMM fabric you need the following:

- **A service template:** VMM uses a service template to automate network controller deployment. Service templates for the network controller support multi-node deployment on generation 1 and generation 2 VMs.
- **A virtual hard disk:** The service template needs a prepared virtual hard disk that's imported into the VMM library. This virtual hard disk is used for network controller VMs.
 - The virtual hard disk must be running Windows Server 2016 with latest patches installed.
 - It can be in VHD or VHDX format.
- **A management logical network:** That models your physical management network's connectivity for the VMM hosts, network controller hosts, and tenant VM hosts.
- **A logical switch:** To provide the management logical network with connectivity to the network controller VMs.
- **An SSL certificate:** To authenticate communications between the VMM server and the network controller.
- **An HNV provider logical network and tenant VM networks:** To validate the network controller deployment.
- **Other prerequisites:** [Verify](#) other requirements.

Deployment steps

Here's what you need to do to set up a SDN network controller

1. **Configure hosts and physical network infrastructure:** You need access to your physical network devices to configure VLANs, routing etc. You also need Hyper-V hosts to host the SDN infrastructure and tenant VMs. [Learn more.](#)
2. **Prepare a virtual hard disk:** You can prepare a virtual hard disk for the network controller service template in VHD or VHDX format, as appropriate for the service template generation you choose.
3. **Download the service templates:** Download the network controller service templates and import them to the VMM library.
4. **Set up Active Directory security groups:** You'll need an Active Directory security group for network

controller management, and another security group for network controller clients. Each group will need at least one user account in it.

5. **Set up a VMM library share.** You can have an optional library file share for keeping diagnostic logs. This library share will be accessed by the network controller to store diagnostics information throughout its lifetime.
6. **Set up a VMM host group:** Set up a dedicated host group for all of the SDN Hyper-V hosts.

NOTE

Hosts must be running Windows Server 2016 with latest patches installed, and have the Hyper-V role enabled.

7. **Create the management logical network:** Create a logical network to mirror management network connectivity for the VMM host, network controller hosts, and tenant VM hosts. If you want to allocate static IP addresses from a pool, create a pool on this logical network.
8. **Create and deploy a management logical switch:** You create the logical switch, and deploy it on network controller hosts, to provide connectivity to the management network for network controller VMs.
9. **Set up a certificate:** You need an SSL certificate for secure/HTTPS communication with the network controller.
10. **Import the template:** Import and customize the network controller service template.
11. **Deploy the service:** Deploy the network controller service using the service template. Then add it as a VMM service.

Prepare a virtual hard disk

1. Prepare the VHD or VHDX based on the type of template you would like to use.
2. After you prepare the hard disk, install the latest Windows Server 2016 updates, and any language packs you need if you have a non-English environment.
3. Import the VHD/VHDX files to the VMM library. [Learn more.](#)

Download the network controller service template

1. Download the SDN folder from the [Microsoft SDN GitHub repository](#) and copy the templates from **VMM > Templates > NC** to a local path on the VMM server.
2. Extract the contents to a folder on a local computer.
3. Refresh the library, You'll import the service templates, later.

NOTE

The custom resource files are used when setting up the network controller, and other SDN components (software load balancer, RAS gateway).

The NC folder contains Four service templates and Five custom resource folders. These are summarized in the following table:

Templates and resource files

NAME	TYPE	DETAILS
Network Controller Production Generation 1 VM.xml	Template	Three-node network controller for generation 1 VMs
Network Controller Production Generation 2 VM.xml	Template	Three-node network controller for generation 2 VMs
Network Controller Standalone Generation 1 VM.xml	Template	Single-node network controller for generation 1 VMs
Network Controller Standalone Generation 2 VM.xml	Template	Single-node network controller for generation 2 VMs
NcSetup.cr	Custom resource file	A library resource containing scripts used to set up the network .
ServerCertificate.cr	Custom resource file	Library resource containing the private key for the network controller in .pfx format.
NcCertificate.cr	Custom resource file	Library resource containing the trusted root certificate (.CER) for the network controller. This is used for secure communications between the network controller and other sub-services (For example, SLB MUXes).
TrustedRootCertificate.cr	Custom resource file	Library resource containing the CA public key (.cer), imported as the trusted root certificate to validate the SSL certificate.
EdgeDeployment.cr	Template	Used for installing SLB MUX roles and gateway roles (for example, VPN)

Set up Active Directory groups

Create security groups for network controller management and clients.

1. In **Active Directory Users and Computers**, create a security group for network controller management.
 - In the group, add all the users who have permission to configure the network controller. For example, create a group named Network Controller Admins.
 - All of the users that you add to this group must also be members of the Domain Users group in Active Directory.
 - The group for network controller management should be a domain local group. Members of this group will be able to create, delete, and update the deployed network controller configuration.
 - Create at least one user account that is a member of this group and have access to its credentials. After the network controller is deployed, VMM can be configured to use the user account credentials to establish communication with the network controller.
2. Create another security group for network controller clients
 - Add users with permission to configure and manage networks using network controller. For example, create a group named Network Controller Users.
 - All of the users that you add to the new group must also be members of the Domain Users group in

Active Directory.

- All Network Controller configuration and management is performed using Representational State Transfer (REST).
- The group should be a Domain Local group. After the network controller is deployed, any members of this group will have permissions to communicate with the network controller via the REST based interface.
- Create at least one user account that is a member of this group. After the network controller is deployed, VMM can be configured to use the user account credentials to establish communication with the network controller.

Create a library share for logging

1. Optionally [create a file share](#) in the VMM library to keep diagnostic logs.
2. Ensure that the share can be accessed by the network controller. The network controller accesses the share to store diagnostic information. Note the credentials for the account that will have write access to the share.

Set up host groups

1. [Create a dedicated host group](#) for Hyper-V hosts that will be managed by SDN.
2. Make sure Hyper-V hosts are running Windows Server 2016 with latest patches installed.

Create the management logical network

You create a management logical network in VMM, to mirror your physical management network.

- The logical network provides network connectivity settings for the VMM host, network controller hosts, and tenant VM hosts.
- We recommend that you create this logical network specifically to provide connectivity for infrastructure VMs that are managed by the network controller.
- If you already have a VMM logical network that's configured with **Create a VM Network with the same name to allow virtual machines to access this logical network directly**, then you can reuse this logical network to provide management connectivity to network controller.

use the following procedure to create management logical network:

1. Click **Fabric > Networking**. Right-click **Logical Networks > Create Logical Network**.
2. Specify a **Name** and optional **Description**.
3. In **Settings** select **One Connected Network**. All management networks need to have routing and connectivity between all hosts in that network. Select **Create a VM network with the same name to allow virtual machines to access this logical network directly** to automatically create a VM network for your management network.
4. Click **Network Site > Add**. Select the host group for the hosts that will be managed by the network controller. Insert your management network IP subnet details. This network should already exist and be configured in your physical switch.
5. Review the **Summary** information and click **Finish** to complete.

Create an IP address pool

If you want to allocate static IP addresses to network controller VMs, create an IP address pool in the management logical network. If you're using DHCP you can skip this step.

1. In the VMM console, right-click the management logical network and select **Create IP Pool**.
2. Provide a **Name** and optional description for the pool and ensure that the management network is selected for the logical network.

3. In **Network Site** panel, select the subnet that this IP address pool will service.
4. In **IP Address range** panel, type the starting and ending IP addresses.
5. To use an IP as REST IP, type one of the IP addresses from the specified range in **IP addresses to be reserved for other uses** box. In case you want to use the REST End Point, skip this step.
 - Don't use the first three IP addresses of your available subnet. For example, if your available subnet is from .1 to .254, start your range at .4 or greater.
 - If the nodes are in the same subnet, you must provide REST IP address. If the nodes are in different subnets, you must provide a REST DNS name.
6. Specify the default gateway address and optionally configure DNS and WINS settings
7. In **Summary** page, review the settings and click **Finish** to complete the wizard.

Create and deploy a management logical switch

You need to deploy a logical switch on the management logical network. The switch provides connectivity between the management logical network and the network controller VMs.

1. In the VMM console, click **Fabric > Networking > Create Logical Switch**. Review the Getting Started information and click **Next**.
2. Provide a **Name** and optional description. Select **No Uplink Team**. If you need teaming, select **Embedded Team**.

NOTE

Do not use **Team**.

3. For minimum bandwidth mode, choose the **Weight** option.
4. In **Extensions**, clear all the switch extensions. This is important. If you select any of the switch extensions at this stage, it could block the network controller onboarding later.
5. You can optionally add a virtual port profile and choose a port classification for host management.
6. Select an existing uplink port profile, or click **Add > New Uplink Port Profile**. Provide a **Name** and optional description. Use the defaults for load balancing algorithm and teaming mode. Select all the network sites in the management logical network.
7. Click **New Network Adapter**. This adds a host virtual network adapter (vNIC) to your logical switch and uplink port profile, so that when you add the logical switch to your hosts, the vNICs get added automatically.
8. Provide a **Name** for the vNIC. Verify that the management VM network is listed in **Connectivity**.
9. Select **This network adapter will be used for host management > Inherit connection settings from the host adapter**. This allows you to take the vNIC adapter settings from the adapter that already exists on the host. If you created a port classification and virtual port profile earlier, you can select it now.
10. In **Summary** review the information and click **Finish** to complete the wizard.

Deploy the logical switch

You must deploy the management logical switch on all of the hosts where you intend to deploy the NC. These hosts must be a part of VMM host group that you created earlier [Learn more](#).

Set up the security certificates

You need an SSL certificate that will be used for secure/HTTPS communication with the network controller. You can use the following methods:

- **Self-signed certificate:** You can generate a self-signed certificate, and export it with the private key protected with a password.
- **Certificate Authority (CA) certificate:** You can use a certificate signed by a CA.

Use a self-signed certificate

The following example creates a new self-signed certificate, and should be run on the VMM server.

NOTE

- You can use an IP address as the DNS name, but this is not recommended as it restricts the network controller to a single subnet.
- You can use any friendly name for the network controller.
- For multi-node deployment, The DNS name should be the REST name you want to use.
- For single-node deployment, the DNS name should be the network controller name followed by the full domain name.

DEPLOYMENT	SYNTAX	EXAMPLE
Multi-node	<pre>New-SelfSignedCertificate - KeyUsageProperty All -Provider "Microsoft Strong Cryptographic Provider" -FriendlyName " <YourNCComputerName>" -DnsName @"("<NCRESTName>")</pre>	<pre>New-SelfSignedCertificate - KeyUsageProperty All -Provider "Microsoft Strong Cryptographic Provider" -FriendlyName "MultiNodeNC" -DnsName @"("NCCluster.Contoso.com")</pre>
Single-node	<pre>New-SelfSignedCertificate - KeyUsageProperty All -Provider "Microsoft Strong Cryptographic Provider" -FriendlyName " <YourNCComputerName>" -DnsName @"("<NCFQDN>")</pre>	<pre>New-SelfSignedCertificate - KeyUsageProperty All -Provider "Microsoft Strong Cryptographic Provider" -FriendlyName "SingleNodeNC" -DnsName @"("SingleNodeNC.Contoso.com")</pre>

Export the self-signed certificate

Export the certificate and its private key in .pfx format.

1. Open the **Certificates** snap-in (certlm.msc) and locate the certificate in Personal/Certificates.
2. Select the certificate > **All Tasks** > **Export**.
3. Select **Yes, export the private key** option, click **Next**.
4. Choose **Personal Information Exchange - PKCS #12 (.PFX)** and accept the default to **Include all certificates in the certification path if possible**.
5. Assign the **Users/Groups** and a password for the certificate you are exporting, click **Next**.
6. On the **File to export** page, browse the location where you want to place the exported file, and give it a name.
7. Similarly, export the certificate in .CER format

NOTE

To export to .CER format, uncheck the **Yes, export the private key** option.

8. Copy the .PFX to the ServerCertificate.cr folder.
9. Copy the .CER file to the NCCertificate.cr folder.

When you are done, refresh these folders and ensure that you have these certificates, copied.

Use a CA

1. Request a CA-signed certificate. For a Windows-based enterprise CA, [request certificates using the certificate request Wizard](#).
2. Make sure that the certificate includes the serverAuth EKU, specified by the OID 1.3.6.1.5.5.7.3.1. In addition, the certificate subject name must match the DNS name of the network controller.
3. Copy the .PFX to the ServerCertificate.cr folder.
4. Copy the .CER file to the NCCertificate.cr folder.
5. Copy the public key of the CA in .CER format to TrustedRootCertificate.cr.

NOTE

Ensure that the enterprise CA is configured for certificate auto enrollment.

Set up the service template

Import the template, and update the parameters for your environment.

Import the template

Import the service template into the VMM library. For this example we'll import the generation 2 template.

1. Click **Library > Import Template**.
2. Browse to your service template folder, select the **Network Controller Production Generation 2 VM.xml** file.
3. Update the parameters for your environment as you import the service template. Review the details and then click **Import**.
 - **WinServer.vhdx** Select the base virtual hard drive image that you prepared earlier.
 - **NCSetup.cr**: Map to the NCSetup.cr library resource in the VMM library.
 - **ServerCertificate.cr**: Map to the ServerCertificate.cr resource in the VMM library. In addition, put the .pfx SSL certificate that you prepared earlier inside this folder. Make sure you only have one certificate in the ServerCertificate.cr folder.
 - **TrustedRootCertificate.cr**: Map to the TrustedRootCertificate.cr folder in your VMM library. If you don't need a trusted root certificate, this resource still needs to be mapped to a CR folder. However, the folder should be left empty.
4. Once done, ensure the Job is complete.

Customize the template

You can customize the service template to meet any specific requirements related to your organization, such as product key, IP assignment, DHCP, MAC Spoofing and High availability. You can also customize properties for objects such as host groups, host clusters, service instances.

As an example, here are the steps to enter the product key, enable DHCP and high availability:

1. In the VMM library, select the service template, and open it in designer mode.
2. Double-click the computer tier to open the Windows Server Network Controller Properties page.
3. To specify a product key, click **OS Configuration > Product Key**, and specify the key shared by CCEP.

4. To enable high availability, click **Hardware configuration** > **Availability**, select the **Make the Virtual machine highly available** check box.
5. To enable dynamic IP configuration and leverage DHCP for network controller management, click network adapter on the designer, and change the IPV4 address type to **Dynamic**.

NOTE

- If you customize the template for high availability, ensure you deploy this on clustered nodes.
- While configuring your Network Controller and specifying FQDN as the REST name, don't pre-create Host A record for your primary NC node in your DNS. This may impact Network Controller connectivity once primary NC node changes. This is applicable even if you are deploying the NC by using the SDN Express or VMM Express script.

Deploy the network controller

1. Select the network controller service template > **Configure Deployment**. Type a service name, and select a destination for the service instance. The destination must map to the dedicated host group containing hosts that will be managed by network controller.
2. Configure the deployment settings as described in the table below.
3. It's normal for the virtual machine instances to be initially red. Click **Refresh Preview** to have the deployment service automatically find suitable hosts for the virtual machines to be created.
4. After you configure these settings, click **Deploy Service** to begin the service deployment job.

NOTE

Deployment times will vary depending on your hardware but are typically between 30 and 60 minutes. Note that if you're not using a volume licensed VHD\VHDX, or if the VHD\VHDX doesn't supply the product key using an answer file, then the deployment stops at the **Product Key** page, during network controller VM provisioning. You need to manually access the VM desktop, and either skip or enter the product key.

5. If the network controller deployment fails, delete the failed service instance, before you retry the network controller deployment. Click **VMs and Services** > **All Hosts** > **Services**, and delete the instance.

Deployment settings

SETTING	REQUIREMENT	DESCRIPTION
ClientSecurityGroup	Required	Name of the security group that you created, containing network controller client accounts.
DiagnosticLogShare	Optional	File share location where the diagnostic logs will be periodically uploaded. If this is not provided, the logs are stored locally on each node.
DiagnosticLogShareUsername	Optional	Full user name (including domain name) for an account that has access permissions to the diagnostic log share. In the format: [domain]\[username].

SETTING	REQUIREMENT	DESCRIPTION
DiagnosticLogSharePassword	Optional	The password for the account specified in the DiagnosticLogShareUsername parameter.
LocalAdmin	Required	<p>Select a Run as account in your environment, which will be used as the local administrator on the network controller virtual machines.</p> <p>Note: while creating Run as accounts, uncheck the validate domain credentials option if you are creating a local account.</p> <p>User name should be .\Administrator (create it if it doesn't exist).</p>
Management	Required	Select the management logical network you created earlier.
MgmtDomainAccount	Required	Select a Run as account in your environment which will be used to prepare the network controller. This user must be a member of the management security group, specified below, which has privileges to manage the network controller.
MgmtDomainAccountName	Required	<p>This must be the full user name (including domain name) of the Run as account mapped to MgmtDomainAccount.</p> <p>The domain user name will be added to the Administrators group during deployment.</p>
MgmtDomainAccountPassword	Required	Password for the management Run as account mapped to MgmtDomainAccount.
MgmtDomainFQDN	Required	FQDN for the Active directory domain that the network controller virtual machines will join.
MgmtSecurityGroup	Required	Name of the security group you created previously containing network controller management accounts.
RestEndPoint	Required	<p>Enter the RESTName you used when preparing the certificates. This parameter isn't used for standalone templates.</p> <p>If the nodes are in the same subnet, you must provide the REST IP address. If the nodes are in different subnets, provide the REST DNS name.</p>

SETTING	REQUIREMENT	DESCRIPTION
ServerCertificatePassword	Required	Password to import the certificate into the machine store.

NOTE

Windows 2019 onwards, the Network Controller machines must be provided permission to register and modify the SPN in the Active Directory. For more details, see [Kerberos with Service Principal Name](#).

Add the network controller service to VMM

After the network controller service is successfully deployed, the next step is to add it to VMM as a network service.

1. In **Fabric**, right-click **Networking** > **Network Service**, and click **Add Network Service**.
2. The **Add Network Service Wizard** starts. Specify a name and optional description.
3. Select **Microsoft** for the manufacturer and for model select **Microsoft network controller**.
4. In **Credentials**, provide the Run As account you want to use to configure the network service. This should be the same account that you included in the network controller clients group.
5. For the **Connection String**:
 - In multi-node deployment, **ServerURL** should use the REST endpoint, and **servicename** should be the name of the network controller instance.
 - In single node deployment, **ServerURL** should be the network controller FQDN and, **servicename** must be the network controller service instance name. Example:


```
serverurl=https://NCCluster.contoso.com;servicename=NC_VMM_RTM
```
6. In **Review Certificates**, a connection is made to the network controller virtual machine to retrieve the certificate. Verify that the certificate shown is the one you expect. Ensure you select **These certificates have been reviewed and can be imported to the trusted certificate storebox**.
7. On the next screen, click **Scan Provider** to connect to your service and list the properties and their status. This is also a good test of whether or not the service was created correctly, and that you're using the right connect string to connect to it. Examine the results, and check that `isNetworkController = true`. When it completes successfully click **Next**.
8. Configure the host group that your network controller will manage.
9. Click **Finish** to complete the wizard. When the service has been added to VMM, it will appear in the **Network Services** list in the VMM console. If the network service isn't added, check **Jobs** in the VMM console to troubleshoot.

Validate the deployment

You can optionally validate the network controller deployment. To do this:

1. Create **HNV provider** network (the backend network), managed by the network controller for tenant VM connectivity. This network is used to validate that the network controller has been deployed successfully, and that tenant VMs within the same virtual network can ping each other. This network should exist in your physical network infrastructure, and all SDN fabric hosts should have physical connectivity to it.
2. After creating the HNV provide network, you configure two tenant VM networks on top of it. Create VM

networks and IP address pools, and then deploy the tenant VMs. You can also test connectivity between two tenant VMs deployed on different hosts to ensure the network controller is deployed correctly.

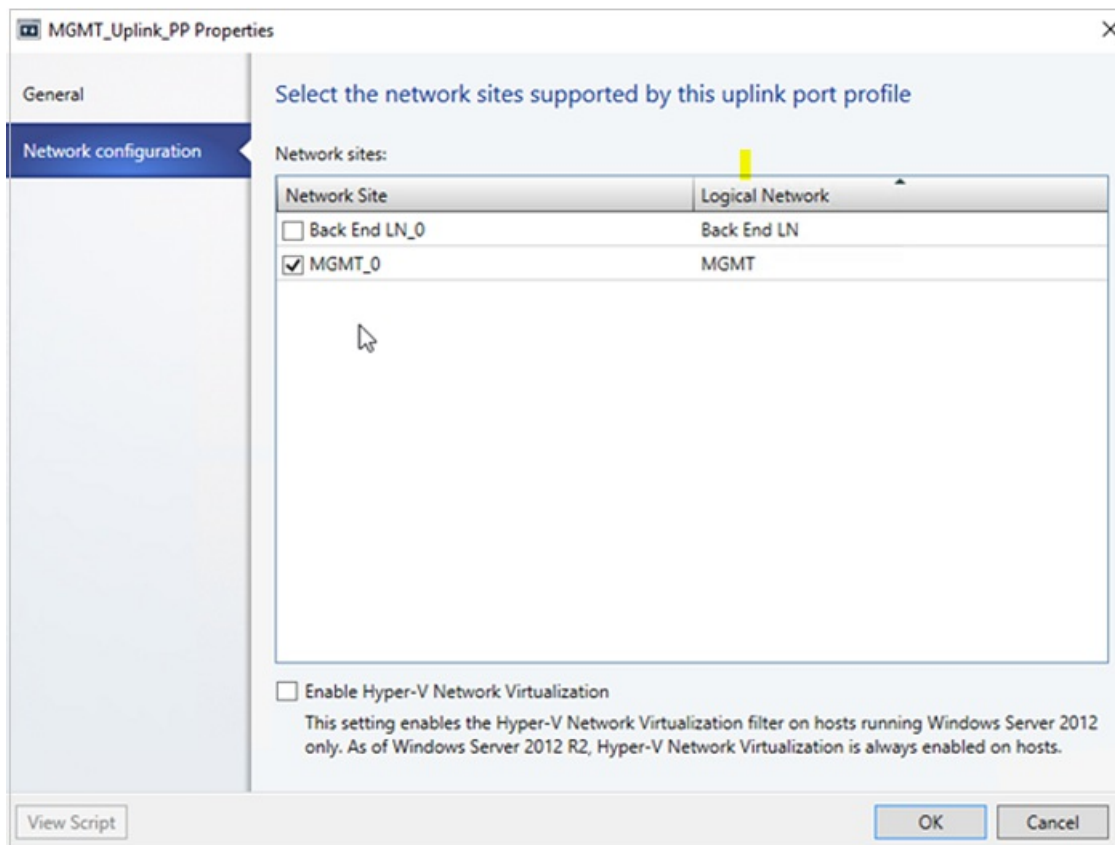
Create the HNV provider network

1. Start the **Create Logical Network Wizard**. Type a name and optional description for this network.
2. In **Settings** verify that **One Connected Network** is selected, since all HNV Provider networks need to have routing and connectivity between all hosts in that network. Ensure you check **Allow new VM networks created on this logical network to use network virtualization**. In addition check **Managed by the network controller**.
3. In **Network Site** add the network site information for your HNV Provider network. This should include the host group, subnet, and VLAN information for the network.
4. Review the **Summary** information and complete the wizard.

Create the IP address pool

The HNV provider logical network needs an IP address pool, even if DHCP is available on this network. If you have more than one subnet on the HNV provider network, create a pool for each subnet.

1. Right-click the HNV Provider logical network > **Create IP Pool**.
2. Provide a name and optional description, and ensure that the HNV Provider logical network is selected for the logical network.
3. In **Network Site** you need to select the subnet that this IP address pool will service. If you have more than one subnet as part of your HNV provider network, you need to create a static IP address pool for each subnet. If you have only one site (for example, like the sample topology) then you can just click **Next**.
4. In **IP Address range** configure the starting and ending IP address. Don't use the first IP address of your available subnet. For example, if your available subnet is from .1 to .254, start your range at .2 or greater.
5. Next, configure the default gateway address. Click **Insert** next to the **Default gateways** box, type the address and use the default metric. Optionally configure DNS and WINS.
6. Review the summary information and click **Finish** to complete the wizard.
7. As part of network controller on-boarding, the switch that you deployed on the hosts for the Management logical network connectivity was converted to an SDN switch. This switch can now be used to deploy a network controller managed network including the HNV provider logical network. Ensure you select the network site corresponding to the HNV provider logical network in the uplink port profile settings for the Management logical switch.



The HNV provider logical network is now accessible to all the hosts in the network controller managed host group.

Create tenant VM networks and IP pools

Now, create two VM networks and IP pools for two tenants in your SDN infrastructure, to test connectivity.

NOTE

- Do not use the first IP address of your available subnet. For example, if your available subnet is from .1 to .254, start your range at .2 or greater.
- Currently you can't create a VM network with **No Isolation** for logical networks that are managed by the network controller. You must choose the **Isolate using Hyper-V Network Virtualization** isolation option when creating VM Networks associated with HNV Provider logical networks.
- Since the network controller is not yet tested with IPv6, use IPv4 for both the logical network and the VM network when you create a VM network.

1. [Create a VM network](#) for each tenant.
2. [Create an IP address pool](#) for each VM network.

Create tenant virtual machines

Now you can create tenant virtual machines connected to the tenant virtual network. Note that:

- Ensure that your tenant virtual machines allow IPv4 ICMP through their firewall. By default, Windows Server blocks this.
- You can run the following PowerShell cmdlet to allow ICMPv4 in through the firewall with **New-NetFirewallRule -DisplayName "Allow ICMPv4-In" -Protocol ICMPv4**

1. If you want to create a VM from an existing hard disk, [follow these instructions](#).
2. After you deploy at least two VMs connected to the network, you can ping one tenant virtual machine from the other tenant virtual machine to validate that the network controller has been deployed as a network service successfully, and that it can manage the HNV Provider network so that tenant virtual machines can

ping each other.

Remove the network controller from the SDN fabric

Use [these steps](#) to remove the network controller from the SDN fabric.

Next steps

[Create a software load balancer](#)

Set up an SDN software load balancer in the VMM fabric

3/14/2019 • 12 minutes to read

This article describes how to deploy a Software Defined Networking (SDN) software load balancer (SLB) in the System Center - Virtual Machine Manager (VMM) fabric.

The SLB enables even distribution of tenant and tenant customer network traffic among virtual network resources, so that multiple servers can host the same workload to provide high availability and scalability. [Learn more.](#)

You can use VMM to deploy a network controller and a software load balancer. After you set up the SLB, you can leverage the multiplexing and NAT capabilities in your SDN infrastructure.

Before you start

Ensure the following:

- **Planning:** Read about planning a software defined network, and review the planning topology in [this](#) document. The diagram shows a sample 4-node setup. The setup is highly available with Three network controller nodes (VM), and Three SLB/MUX nodes. It shows Two tenants with One virtual network broken into Two virtual subnets to simulate a web tier and a database tier. Both the infrastructure and tenant virtual machines can be redistributed across any physical host.
- **Network controller:** You should have an [SDN network controller](#) deployed in the VMM fabric, so that you have the compute and network infrastructure running before you set up the load balancing.
- **SSL certificate:** To import the SLB service template you'll need to prepare an SSL certificate. You made the certificate available during network controller deployment. To use the certificate you prepared in network controller deployment for SLB, right-click the certificate and export it without a password in .CER format. Place it in the library, in the NCCertificate.CR folder you created when you set up the network controller.
- **Service template:** VMM uses a service template to automate SLB deployment. Service templates support multi-node deployment on generation 1 and generation 2 VMs.
- **SLB VMs:** All the SLB virtual machines must be running Windows Server 2016 with the latest patches installed.
- **HNV Network:** Ensure that you created the Provider HNV network as part of NC validation. [Learn more.](#)

Deployment steps

1. **Prepare the SSL certificate:** Put the certificate in the VMM library.
2. **Download the service template:** Download the service template that you need to deploy the SLB/MUX.
3. **Create the transit logical networks:** You need to create logical networks:
 - A logical network to mirror the transit (Frontend) physical network.
 - Private virtual IP (VIP) and public VIP networks, to assign VIPs to the SLB service.
4. **Create private and public VIP logical networks:** Private virtual IP (VIP) and public VIP networks, to assign VIPs to the SLB service.
5. **Import the service template:** Import and customize the SLB service template.
6. **Deploy SLB:** Deploy SLB as a VMM service, and configure the service properties.
7. **Validate the deployment:** Configure BGP peering between the SLB/MUX instance and a BGP router, assign a public IP address to a tenant VM or service, and access the VM or service from outside the network.

Prepare the certificate

Ensure that the SSL certificate that you created during the NC deployment is copied to NCCertificate.CR folder.

Download the service template

1. Download the SDN folder from the [Microsoft SDN GitHub repository](#) and copy the templates from **VMM** > **Templates** > **SLB** to a local path on the VMM server.
2. Extract the contents to a folder on a local computer. You'll import them to the library later.

The download contains Two templates:

- The SLB Production Generation 1 VM.xml template is for deploying the SLB Service on generation 1 virtual machines.
- The SLB Production Generation 2 VM.xml is for deploying the SLB Service on Generation 2 virtual machines.

Both the templates have a default count of three virtual machines which can be changed in the service template designer.

Create the transit logical network

1. Open the **Create logical network Wizard**, and type a **Name** and optional description.
2. In **Settings**, select **One Connected Network**. Select **Create a VM network with the same name** box to allow virtual machines to access this logical network directly, and **Managed by the network controller**.
3. In **Network Site**, add the network site information for your subnet.
4. Review the **Summary** information and complete the logical network wizard.

Create an IP address pool for the transit logical network

This is the IP address pool where DIPs are assigned to the SLB/MUX virtual machines and BGP Peer virtual machine (if deployed).

Note:

- Ensure you use the IP address range that corresponds to your transit network IP address space. Don't include the first IP address of your subnet in the IP pool you are about to create. For example, if your available subnet is from .1 to .254, start your range at .2.
- After you create the Transit logical network, ensure you associate this logical network with the Management switch uplink port profile you created during the network controller deployment.

Create the IP address pool:

1. Right-click the logical network > **Create IP Pool**.
2. Provide a **Name** and optional description for the IP Pool and ensure that the correct logical network is selected.
3. In **Network Site**, select the subnet that this IP address pool will service. If you have more than one subnet as part of your HNV provider network, you need to create a static IP address pool for each subnet. If you have only one site (for example, like the sample topology) then you can just click **Next**.
4. In **IP Address range** configure the starting and ending IP address. Don't use the first three IP addresses of your available subnet. For example, if your available subnet is from .1 to .254, start your range at .4 or greater.
5. Next, configure the default gateway address. Click **Insert** next to the **Default gateways** box, type the address and use the default metric. Optionally configure DNS and WINS.

6. Review the summary information and click **Finish** to complete the wizard.

Note: Ensure you associate the logical network with the management switch uplink port profile.

Create private and public VIP logical networks

You need a private VIP address pool to assign a VIP, and a public VIP, to the SLB Manager service. Note that the procedure for creating both is similar but there are some differences.

Create a private VIP:

1. Start the **Create logical network Wizard**. Type a **Name** and optional description for this network.
2. In **Settings**, select **One Connected Network**. Select **Create a VM network with the same name to allow virtual machines to access this logical network directly**. Select **Managed by the network controller**.
3. In **Network Site**, add the network site information for your private VIP logical network.
4. Review the **Summary** information and complete the wizard.

Create a public VIP:

1. Start the **Create logical network Wizard**. Type a **Name** and optional description for this network.
2. In **Settings**, select **One Connected Network**. Select **Create a VM network with the same name to allow virtual machines to access this logical network directly**. Select **Managed by the network controller** and **Public IP Address Network**.
3. In **Network Site**, add the network site information for your public VIP logical network.
4. Review the **Summary** information and complete the wizard.

Create IP address pools for the private and public VIP networks

1. Right-click the private VIP logical network > **Create IP Pool**.
2. Provide a **Name** and optional description for the IP Pool and ensure that the correct logical network is selected.
3. Accept the default network site, and click **Next**.
4. In **IP Address range** configure the starting and ending IP address.

Note: Don't use the first IP address of your available subnet. For example, if your available subnet is from .1 to .254, start your range at .2 or greater.

5. In **IP addresses reserved for load balancer VIPs**, type the IP address range in the subnet. It should match the start and end addresses you specified.
6. You don't need to provide gateway, DNS or WINS information because this pool is used to allocate IP addresses for VIPs through the network controller only. Click **Next** to skip these screens. type the address and use the default metric. Optionally configure DNS and WINS.
7. Review the summary information and click **Finish** to complete the wizard.
8. Repeat the procedure for the public VIP logical network, this time type the IP address range for the public network.

Import the service template

Import the service template into the VMM library. For this example, we'll import the generation 2 template.

1. Click **Library > Import Template**.
2. Browse to your service template folder, select the **SLB Production Generation 2 VM.xml** file.

3. Update the parameters for your environment as you import the service template. Note that the library resources were imported during network controller deployment.
 - **WinServer.vhdx**: Select the virtual hard drive image that you prepared and imported earlier, during the network controller deployment.
 - **NCCertificate.CR**: This library resource contains scripts used to set up the network controller. Map to the NCCertificate.cr library resource in the VMM library.
 - **EdgeDeployment.CR**: Map to the EdgeDeployment.cr library resource in the VMM library.
4. Remember that you should have copied the .CER certificate that you previously created to the **NCCertificate.CR** folder.
5. On the **Summary** page, review the details and click **Import**.

Note: You can customize the service template. [Learn more](#).

Deploy the SLB service

Now deploy an SLB/MUX service instance.

1. Select the **SLB Production Generation 2 VM.xml** service template > **Configure Deployment**. Type a **Name** and optional destination for the service instance. The destination must map to a host group that contains the hosts you've configured.
2. In the **Network Settings** section, map **TransitNetwork** to your transit VM network and **ManagementNetwork** to your management VM network.

Note: The **Deploy Service** dialog appears after mapping is complete. It is normal for the virtual machine instances to be initially Red. Click **Refresh Preview** to automatically find suitable hosts for the virtual machine.

3. On the left of the **Configure Deployment** window, configure the settings as detailed in the following table:

SETTING	REQUIREMENT	DESCRIPTION
Transit network	Required	Your transit VM network.
LocalAdmin	Required	Select a Run As Account in your environment, which will be used as the local Administrator on the virtual machines. The user name should be Administrator.
Management network	Required	Choose the management VM network that you created for host management.
MgmtDomainAccount	Required	Select a Run As Account with permissions to add the SLB/MUX virtual machines to the Active Directory domain associated with the network controller. This can be the same account you used in MgmtDomainAccount while deploying the network controller.

SETTING	REQUIREMENT	DESCRIPTION
MgmtDomainFQDN	Required	FQDN for the Active directory domain that the SLB/MUX virtual machines will join.
SelfSignedConfiguration	Required	Specify True if the certificate you are using is self-signed.

- After you configure these settings, click **Deploy Service** to begin the service deployment job. Deployment times will vary depending on your hardware but are typically between 30 and 60 minutes.
- If you are not using a volume licensed VHDX, or if the VHDX doesn't have the product key from an answer file, then deployment will stop at the **Product Key** page, during SLB/MUX VM provisioning. You need to manually access the VM desktop, and either skip or enter the product key.
- When the service deployment job is complete, verify that your service appears in **VMs and Services > Services > VM Network Information for Services**. Right-click the service and verify that the state is **Deployed** in **Properties**.

After deployment, verify that the service appears in **All Hosts > Services > VM Network Information for Services**. Right-click the SLB MUX service > **Properties**, and verify that the state is **Deployed**. If the SLB/MUX deployment fails, ensure you delete the failed service instance, before you try to deploy the SLB, once again.

If you want to scale-in or scale-out a deployed software load balancer service instance, [read this blog](#).

NOTE

After the SLB service is deployed, disable DNS registration on the virtual network adapter, connected to the transit VM network on all the SLB MUX VMs.

Configure the SLB role and SLB/MUX properties

Note: Before you proceed, ensure you created the HNV PA Logical Network.

Now that the service is deployed, you can configure its properties. you'll need to associate the SLB service instance that you deployed with network controller, and then configuring BGP peering between the SLB/MUX instance and a TOR switch or a BGP router peer.

- Click **Fabric > Network Service**. Right-click the **network controller** service > **Properties**.
- Select the **Services** tab > **Load Balancer Role > Associated Service > Browse**.
- Select the SLB/MUX service instance you created earlier. Select a Run As Account.
- For the **Management IP address**, use an IP address from the private VIP pool you created earlier. Optionally specify the IP address ranges to be excluded from the outbound NAT. Under **SLBM VIP Pools**, select both the private and public VIP pools for publishing to NC.
- Click the SLB/MUX instance listed under **Load Balancer Role** in the wizard. Type the local ASN for your datacenter and details for the devices or BGP peers the SLB/MUX can peer with.
- Click **OK**.

The SLB service instance is now associated with the SLBM service, and you should see the SLB/MUX virtual machine instance with all the settings listed under the **Load Balancer role**.

Validate the deployment

After you deploy the SLB/MUX, you can validate the deployment by configuring BGP peering between the

SLB/MUX instance and a BGP router, assigning a public IP address to a tenant virtual machine or Service, and accessing the tenant virtual machine or service from outside the network.

Use the following procedure to validate:

1. Enter your external router details in the wizard. For example:

The screenshot shows the 'Services' configuration window for 'muxvm001.cdm.lab'. On the left is a sidebar with navigation options: Name, Manufacturer and Model, Credentials, Connection String, Host Group, Logical Network Affinity, Extensions, and Services (which is highlighted). The main area is titled 'Services' and contains a 'Load balancer role' section with '1 active nodes' and a 'Gateway manager role' section. The 'BGP settings' section is expanded, showing 'Local ASN: 3159' and a prompt to 'Provide a list of devices to peer with.' Below this is a table with columns 'Name', 'IP Address', and 'ASN'. One row is populated with 'DC-BGP', '193.68.110.13', and '1677'. To the right of the table are 'Add' and 'Remove' buttons. At the bottom of the window are 'View Script', 'OK', and 'Cancel' buttons.

Name	IP Address	ASN
DC-BGP	193.68.110.13	1677

2. Click **OK** to complete the SLB/MUX service instance configuration.
3. Check the **Jobs** window to verify that the **Update Fabric Role with required configuration**, and **Associate service instance with fabric role** jobs have completed successfully.
4. To complete the BGP peering operation, you need to configure BGP to peer with your SLB/MUX instance on the router. If you use a hardware router, you need to consult your vendor's documentation regarding how to setup BGP peering for that device.

You also need to know the IP address of the SLB/MUX instance that you deployed earlier. To do this, you can either log on to the SLB MUX virtual machine and run **ipconfig /all** from the command prompt, or you can get the IP address from the VMM console.

Note: Enter an IP from the transit network.

5. If you create a new VIP pool after peering is complete, you need to advertise all the VIP address pools using the VMM console.

After you validate, you can start using the SLB for load balancing. For related information, see [load balance network traffic](#) and [configure NAT rules](#).

Remove the software load balancer from the SDN fabric

Use [these steps](#) to remove the SLB from the SDN fabric.

Next steps

[Create a RAS gateway](#)

Set up an SDN RAS gateway in the VMM fabric

3/14/2019 • 8 minutes to read

This article describes how to set up a Software Defined Networking (SDN) RAS gateway in the System Center - Virtual Machine Manager (VMM) fabric.

An SDN RAS gateway is a data path element in SDN that enables site-to-site connectivity between two autonomous systems. Specifically, a RAS gateway enables site-to-site connectivity between remote tenant networks and your datacenter using IPsec, Generic Routing Encapsulation (GRE) or Layer 3 Forwarding. [Learn more](#).

Before you start

Ensure the following before you start:

- **Planning:** Read about planning a software defined network, and review the planning topology in [this](#) document. The diagram shows a sample 4-node setup. The setup is highly available with Three network controller nodes (VM), and Three SLB/MUX nodes. It shows Two tenants with One virtual network broken into Two virtual subnets to simulate a web tier and a database tier. Both the infrastructure and tenant virtual machines can be redistributed across any physical host.
- **Network controller:** You should deploy the network controller before you deploy the RAS gateway.
- **SLB:** To ensure that dependencies are handled correctly, you should also deploy the SLB before setting up the gateway. If an SLB and a gateway are configured, you can use and validate an IPsec connection.
- **Service template:** VMM uses a service template to automate GW deployment. Service templates support multi-node deployment on generation 1 and generation 2 VMs.

Deployment steps

To set up a RAS gateway, do the following:

1. **Download the service template:** Download the service template that you need to deploy the GW.
2. **Create the VIP logical network:** Create a GRE VIP logical network. It needs an IP address pool for private VIPs, and to assign VIPs to GRE endpoints. The network exists to define VIPs that are assigned to gateway VMs running on the SDN fabric for a site-to-site GRE connection.
3. **Import the service template:** Import the RAS gateway service template.
4. **Deploy the gateway:** Deploy a gateway service instance, and configure its properties.
5. **Validate the deployment:** Configure site-to-site GRE, IPsec, or L3, and validate the deployment.

Download the service template

1. Download the SDN folder from the [Microsoft SDN GitHub repository](#) and copy the templates from **VMM > Templates > GW** to a local path on the VMM server.
2. Extract the contents to a folder on a local computer. You'll import them to the library later.

The download contains Two templates:

- The EdgeServiceTemplate_Generation 1 VM.xml template is for deploying the GW Service on generation 1 virtual machines.
- The EdgeServiceTemplate_Generation 2 VM.xml is for deploying the GW Service on Generation 2 virtual machines.

Both the templates have a default count of three virtual machines which can be changed in the service template designer.

Create the GRE VIP logical network

1. In the VMM console, run the Create Logical Network Wizard. Type a **Name**, optionally provide a description, and click **Next**.
2. In **Settings**, select **One Connected Network**. Optionally you can select **Create a VM network with the same name**. This setting allows VMs to access this logical network directly. Select **Managed by the Network Controller**, and click **Next**.
3. In **Network Site**, specify the settings:

Here are the sample values:

- Network name: GRE VIP
- Subnet: 31.30.30.0
- Mask: 24
- VLAN ID on trunk: NA
- Gateway: 31.30.30.1

4. In **Summary**, review the settings and finish the wizard.

Create an IP address pool for GRE VIP addresses

1. Right-click the GRE VIP logical network > **Create IP Pool**.
2. Type a **Name** and optional description for the pool, and check that the VIP network is selected. Click **Next**.
3. Accept the default network site and click **Next**.
4. Choose a starting and ending IP address for your range.

Note: Start the range on the second address of your available subnet. For example, if your available subnet is from .1 to .254, start the range at .2.

5. In the **IP addresses reserved for load balancer VIPs** box, type the IP addresses range in the subnet. This should match the range you used for starting and ending IP addresses.
6. You don't need to provide gateway, DNS or WINS information as this pool is used to allocate IP addresses for VIPs through the network controller only. Click **Next** to skip these screens.
7. In **Summary**, review the settings and finish the wizard.

Import the service template

1. Click **Library > Import Template**.
2. Browse to your service template folder. As an example, select the **EdgeServiceTemplate Generation 2.xml** file.
3. Update the parameters for your environment as you import the service template. Note that the library resources were imported during network controller deployment.
 - **WinServer.vhdx**: Select the virtual hard drive image that you prepared and imported earlier, during the network controller deployment.
 - **EdgeDeployment.CR**: Map to the EdgeDeployment.cr library resource in the VMM library.
4. On the **Summary** page, review the details and click **Import**.

Note: You can customize the service template. [Learn more.](#)

Deploy the gateway service

This example uses the generation 2 template.

1. Select the **EdgeServiceTemplate Generation2.xml** service template, and click **Configure Deployment**.
2. Type a **Name** and choose a destination for the service instance. The destination must map to a host group that contains the hosts configured previously for gateway deployment.
3. In **Network Settings**, map the management network to the management VM network.

Note: The **Deploy Service** dialog appears after mapping is complete. It's normal for the VM instances to be initially Red. Click **Refresh Preview** to automatically find suitable hosts for the VM.

4. On the left of the **Configure Deployment** window, configure the following settings:
 - **AdminAccount.** Required. Select a RunAs account that will be used as the local administrator on the gateway VMs.
 - **Management Network.** Required. Choose the Management VM network that you created for host management.
 - **Management Account.** Required. Select a Run as account with permissions to add the gateway to the Active Directory domain associated with the network controller. This can be the same account used for MgmtDomainAccount while deploying the network controller.
 - **FQDN.** Required. FQDN for the Active directory domain for the gateway.
5. Click **Deploy Service** to begin the service deployment job.

Note:

- Deployment times will vary depending on your hardware but are typically between 30 and 60 minutes. If gateway deployment fails, delete the failed service instance in **All Hosts > Services** before you retry the deployment.
- If you aren't using a volume licensed VHDX (or the product key isn't supplied using an answer file), then deployment will stop at the **Product Key** page during VM provisioning. You need to manually access the VM desktop, and either enter the key, or skip it.
- If you want to scale-in or scale-out a deployed SLB instance, read this [blog](#).

Gateway limits

The following are the default limits for NC managed gateway:

- **MaxVMNetworksSupported**= 50
- **MaxVPNConnectionsPerVMNetwork**= 10
- **MaxVMSubnetsSupported**= 550
- **MaxVPNConnectionsSupported**= 250

Override the gateway limits

To override the default limits, append the override string to the network controller service connection string and update in VMM.

- **MaxVMNetworksSupported**= followed by the number of VM networks that can be used with this gateway.
- **MaxVPNConnectionsPerVMNetwork**= followed by the number of VPN Connections that can be created per VM network with this gateway.
- **MaxVMSubnetsSupported**= followed by the number of VM network subnets that can be used with this

gateway.

- **MaxVPNConnectionsSupported**= followed by the number of VPN Connections that can be used with this gateway.

Example:

To override the maximum number of VM networks that can be used with the gateway to 100, update the connection string as follows:

```
serverurl=https://NCCluster.contoso.com;servicename=NC_VMM_RTM; MaxVMNetworksSupported==100
```

Configure the gateway manager role

Now that the gateway service is deployed, you can configure the properties, and associate it with the network controller service.

1. Click **Fabric** > **Network Service** to display the list of network services installed. Right-click the network controller service > **Properties**.
2. Click the **Services** tab, and select the **Gateway Manager Role**.
3. Find the **Associated Service** field under **Service information**, and click **Browse**. Select the gateway service instance you created earlier, and click **OK**.
4. Select the **Run As account** that will be used by network controller to access the gateway virtual machines.

Note: The Run as account must have Administrator privileges on the gateway VMs.

5. In **GRE VIP subnet**, select the VIP subnet that you created previously.
6. In **Public IPv4 pool**, select the pool you configured during SLB deployment. In **Public IPv4 address**, provide an IP address from the previous pool, and ensure you don't select the initial three IP addresses from the range.
7. In **Gateway Capacity**, configure the capacity settings.

The gateway capacity (Mbps) denotes the normal TCP bandwidth that is expected out of the gateway VM. You must set this parameter based on the underlying network speed you use.

IPsec tunnel bandwidth is limited to (3/20) of the gateway capacity. Which means, if the gateway capacity is set to 1000 Mbps, the equivalent IPsec tunnel capacity would be limited to 150 Mbps.

NOTE

The bandwidth limit is the total value of inbound bandwidth and outbound bandwidth.

The equivalent ratios For GRE, and L3 tunnels are 1/5 and 1/2 respectively.

8. Configure the number of reserved nodes for back-up in **Nodes for reserved for failures field**.
9. To configure individual gateway VMs, click each VM and select the IPv4 frontend subnet, specify the local ASN, and optionally add the peering device information for the BGP peer.

Note: You must configure the gateway BGP peers, if you plan to use GRE connections.

The service instance you deployed is now associated with the gateway Manager role. You should see the gateway VM instance listed under it.

Validate the deployment

After you deploy the gateway, you can configure S2S GRE, S2S IPSec, or L3 connection types, and validate them. For additional information, see the following contents:

- [Create and validate site-to-site IPSec connections](#)
- [Create and validate site-to-site GRE connections](#)
- [Create and validate L3 connections](#)

For more information on connection types, see [this](#) article.

Set up the traffic selector from PowerShell

Here is the procedure to setup the traffic selector by using the VMM PowerShell.

1. Create the traffic selector by using the following parameters.

Note: Values used are examples only.

```
$t= new-object Microsoft.VirtualManager.Remoting.TrafficSelector

$t.Type=7 // IPV4=7, IPV6=8

$t.ProtocolId=6 // TCP =6, reference: https://en.wikipedia.org/wiki/List_of_IP_protocol_numbers

$t.PortEnd=5090

$t.PortStart=5080

$t.IpAddressStart=10.100.101.10

$t.IpAddressEnd=10.100.101.100
```

2. Configure the above traffic selector by using **-LocalTrafficSelectors** parameter of **Add-SCVPNConnection** or **Set-SCVPNConnection**.

Remove the gateway from the SDN fabric

Use [these steps](#) to remove the gateway from the SDN fabric.

Set up Software Defined Network (SDN) components in the VMM fabric using PowerShell

12/7/2018 • 3 minutes to read

System Center - Virtual Machine Manager (VMM) can be used to deploy and manage a Software Defined Network (SDN) infrastructure.

You can deploy SDN components in the VMM fabric, including:

- **Network Controller:** The network controller allows you to automate configuration of your network infrastructure, instead of manually configuring network devices and services.
- **RAS Gateway for SDN:** RAS Gateway is a software-based, multitenant, BGP capable router in Windows Server 2016 that is designed for CSPs and Enterprises that host multiple tenant virtual networks using HNV.
- **Software Load Balancing (SLB) for SDN:** SDN in Windows Server 2016 can use Software Load Balancing (SLB) to evenly distribute tenant and tenant customer network traffic among virtual network resources. The Windows Server SLB enables multiple servers to host the same workload, providing high availability and scalability.

There are a couple of ways to deploy these components:

- **VMM console:** Deploy the [network controller](#), [SLB](#), and [RAS gateway](#) manually in the VMM console.
- **PowerShell:** Deploy all components using PowerShell scripts.

Advantages of PowerShell deployment

- Deploy all SDN components with PowerShell scripts.
- Using a script can reduce the introduction of manual errors, and save significant deployment time.
- If you deploy using the script, afterwards you can modify settings in the VMM console, just as you would if you deploy the SDN components manually.
- Like the manual deployment, you have the option of setting up a new management logical network and switch, or to reuse an existing network and switch.
- If the script deployment fails, all changed settings are rolled back, so that you can start again.
- You can turn off deployment for specific components, For example, if you already have network controller deployed, you can deploy SLB and RAS gateway only.

Before you start

- SET-enabled switch deployment isn't currently supported in a PowerShell deployment. You need to deploy the SET-enabled switch out-of-band, and then specify the name of the switch during deployment.
- Check you have the prerequisites for SDN component deployment in place:
 - [Network controller prerequisites](#)
 - [SLB prerequisites](#)
 - [RAS gateway prerequisites](#)

Deployment steps

Here's what you need to do to set up SDN components in VMM with PowerShell.

1. **Configure hosts and physical network infrastructure:** You need access to your physical network devices

to configure VLANs, routing etc. You also need Hyper-V hosts to host the SDN infrastructure and tenant VMs. [Learn more](#).

2. [Prepared virtual hard disk](#) for the service templates in VHD or VHDX format.
3. Download the [network controller](#) service template, the [SLB service](#) template, and the [RAS gateway](#) service template.
4. Import the [network controller](#), [SLB](#), and [RAS gateway](#) templates into the VMM library.
5. [Set up Active Directory security groups](#). One for network controller management, and another for network controller clients. Each group will need at least one user account in it.
6. [Set up a VMM library share](#). You can have an optional library file share for keeping diagnostic logs. This library share will be accessed by the network controller to store diagnostics information throughout its lifetime.
7. [Set up a dedicated VMM host group](#) for all SDN Hyper-V hosts. Note that hosts must be running the latest version of Windows Server 2016, and have the Hyper-V role enabled.
8. [Set up a certificate](#). You need an SSL certificate for HTTPS communications between VMM and the network controller.
9. [Download](#) and run the SDN scripts. There are three scripts:
 - **VMMExpress.ps1**: This script deploys the SDN stack. After you download it, you can your own customizations.
 - **Fabricconfig.psd1**: This file accepts all the inputs for setting up SDN.
 - **Fabricconfig_Example.psd1**: A sample file that contains dummy parameters. You can replace those with your own parameters.

Next steps

[Configure hosts and physical network infrastructure for SDN](#) .

Set up VM networks in SDN using VMM

5/30/2019 • 3 minutes to read

This article provides information about how to create VM networks in an SDN, using System Center - Virtual Machine Manager (VMM).

VM networks are abstract objects that act as an interface to logical networks. In a virtualized network environment, by using the VM networks, you can abstract virtual machines from the underlying logical network.

A logical network can have one or more VM networks associated with it based on its isolation settings.

The following two types of isolation settings are supported in SDN fabric:

- **Network virtualization** : If a logical network is isolated using network virtualization, you can create multiple VM networks for this logical network. Within a VM network, tenants can use any IP addresses regardless of the IP addresses used on other VM networks. As a service provider, you can host workloads from multiple tenants on a single logical network. Tenants can also configure network connections on these VM networks.
- **No Isolation** : If a logical network has no isolation, then, only a single VM network can be associated with it. As a service provider, you can host infrastructure workloads using this type of isolation settings.

NOTE

VLAN isolation is not supported in SDN fabric.

Before you start

Ensure the following:

- Network controller is deployed in the SDN fabric. [Learn more](#).
- Created a logical network with appropriate isolation settings and is set to be managed by the Microsoft network controller. Also, created the IP pools for this logical network.

NOTE

If you want to deploy the VMs with dynamic IP on "no isolation" network, then IP pools are not required.

- By default, VMs connected to a VM network with network virtualization isolation setting, cannot connect to other networks. If you want your VM network to connect to other networks, you need to first deploy [SDN SLB](#) and [SDN gateway](#).

Create a VM network (network virtualization)

1. In the VMM fabric, click **VMs and Services > VM Networks > Create VM Network**.
2. In **Create VM Network Wizard > Name**, type a name, optional description and select a logical network that was created with network virtualization isolation settings.
3. In **Isolation**, select **Isolate using Hyper-V network virtualization**, and select IPv4 for IP address protocols. Click **Next**.

4. In **VM Subnets** click **Add**, specify the name and subnets for VM network and then click **Next**.

NOTE

You can add multiple subnets.

5. In **Connectivity** panel, select the type of connectivity you want to use for this VM network.

NOTE

By default, all virtual machines in a VM network communicate with each other. If you want virtual machines on this VM network to communicate with other networks, configure the following settings in the **Connectivity** page:

- **Connect to another network through a VPN tunnel:** Select this option if you want the virtual machines on this VM network to communicate with other networks, over a VPN. To automatically learn routes between the sites connected through the VPN tunnel, select the **Enable the border gateway protocol** option. Select the **VPN gateway device** that you want to use and confirm the settings.

Based on your selection, the **VPN Connections** and **Border Gateway Protocol** pages appear, complete the settings based on the information provided by the VPN admin.

- **Connect directly to an additional logical network:** Select this option if you want the virtual machines on this VM network to connect directly to an additional logical network. To enable access to internet resources, select **Network Address Translation (NAT)** or select **Direct Routing** to bridge a virtualized IP address space with a physical IP address space.

6. In **Summary**, verify the settings and click **Finish**.

Once the job is successfully completed, you can view the newly created VM network under **VMs and Services > VM Networks**.

NOTE

After you create a VM network with network virtualization, ensure you [create an IP Pool](#) for this VM network.

Create a VM network (no isolation)

NOTE

While creating the logical network, If you have chosen the option **Create VM network with same name to allow virtual machines to access this logical network directly** , then you can skip the following steps.

1. Click **VMs and Services > VM Networks . Create VM Network**.
2. In **Create VM Network Wizard > Name**, enter a name, optional description. Select a One connected logical network for this VM network. Click **Next**.
3. In **Summary** , verify the settings and click **Finish**.

Once the job is successfully completed, you can view the newly created VM network under **VMs and Services > VM Networks**.

NOTE

If you had created an IP pool on the logical network, the same will be directly available for the VM network.

Next steps

[Create an IP pool for a VM network](#)

Configure encrypted networks in SDN using VMM

8/29/2019 • 2 minutes to read

This article explains about how to encrypt the VM networks in software defined network (SDN) using System Center - Virtual Machine Manager (VMM).

Today, network traffic can be encrypted by the guest OS or an application using technologies like IPSec and TLS. However, these technologies are difficult to implement because of their inherent complexity and challenges related to interoperability between systems because of the nature of implementation.

Using the encrypted networks feature in VMM, end-to-end encryption can be easily configured on the VM networks by using the Network Controller (NC). This encryption prevents traffic between two VMs on the same VM network and same subnet, from being read and manipulated. VMM 1801 and later supports this feature.

The control of encryption is at the subnet level and encryption can be enabled/disabled for each subnet of the VM network.

This feature is managed through the SDN Network Controller (NC). If you do not already have a Software Defined Network (SDN) infrastructure with an NC, see [deploy SDN](#) for more information.

NOTE

This feature currently provides protection from third-party and network admins, and doesn't offer any protection against fabric admins. Protection against fabric admins is in the pipeline and will be available soon.

Before you start

Ensure the following prerequisites are met:

- At least two hosts for tenant VMs, to validate the encryption.
- HNV based VM network with encryption enabled and a certificate, which can be created and distributed by fabric administrator. Not that The certificate along with its private key must be stored in the local certificate store of all the hosts, where the VMs (of that network) reside.

Procedure - configure encrypted networks

Use the following steps:

1. Create a certificate and then place the certificate in the local certificate store of all the hosts, where you plan to place the tenant VMs for this validation.

You can either create a self-signed certificate or get a certificate from a CA. For information on how to generate a self-signed certificate and place them in the appropriate locations of each host you will be using, [see this article](#).

NOTE

Make a note of the "Thumbprint" of the certificate that you generate. In the article pointed above in step 2- you do not have to perform the actions detailed in "Creating a Certificate Credential" and "Configuring a Virtual Network for Encryption". You will configure those settings using VMM, in the following steps:

2. Set up an HNV provider network for tenant VM connectivity, which will be managed by the NC. [Learn more](#).
3. Create a tenant VM Network and a subnet. While creating the subnet, select **Enable Encryption** under **VM Subnets**. [Learn more](#).

In the next step, paste the thumbprint of the certificate that you created.

Name

VM Subnets

Connectivity

VPN Connections

Access

Specify VM subnets

Add Remove

Sub1
192.168.20.0/24

VM subnet

Give a name to the IP subnet to be virtualized by the VM Network. Enter IP subnets using CIDR notation, for example: 192.168.1.0/24, FD4A:29CD:184F:3A2C::/64.

Name:
Sub1

Subnet:
192.168.20.0/24

☒ Enable encryption

View Script

OK Cancel

Name

VM Subnets

Encryption

Connectivity

VPN Connections

Access

Encrypted network settings

Certificate thumbprint:

View Script

OK Cancel

4. Create two VMs on two separate physical hosts, and connect them to the above subnet. [Learn more](#).

5. Attach any packet sniffing application on the two network interfaces of the two hosts, where the tenant VMs are placed.
6. Send traffic, ping, HTTP or any other packets, between the two hosts and check the packets in the packet sniffing application. The packets should not have any discernible plaintext like parameters of an HTTP request.

Allow and block VM traffic using SDN port ACLs

5/30/2019 • 3 minutes to read

In System Center Virtual Machine Manager (VMM), you can centrally configure and manage software defined network (SDN) port access control lists (ACLs).

- A port ACL is a set of port ACL rules that filter the traffic at layer 2 port level.
- A port ACL in VMM filters access to a specific VMM network object.
- Each VMM network object can have only one port ACL attached.
- An ACL contains rules, and can be attached to any number of VMM network objects. You can create an ACL without rules, and add the rules later.
- If an ACL has multiple rules, they are applied based on the priority. After a rule matches the criteria and is applied, no other rules are processed.
- SDN Port ACLs can be applied to virtual subnets and virtual network adapters.

NOTE

Port ACL settings are exposed only through PowerShell cmdlets in VMM, and can't be configured in the VMM console.

Using VMM PowerShell, you can also configure Hyper-V port ACLs. For more information, see [Hyper-v port ACLs](#).

This article provides information about how to create and manage SDN port ACLs by using the VMM PowerShell cmdlets.

Before you start

Ensure that [SDN network controller](#) is deployed.

Create a port ACL

1. Open PowerShell in VMM.
2. Create a port ACL.

```
PS C:\> New-SCPortACL -Name "RDPAccess" -Description "PortACL to control RDP access" -ManagedByNC
```

NOTE

The parameter **-ManagedByNC** ensures that the port ACL is managed by Network Controller (NC) and can only be attached to NC managed objects. The cmdlets provided here use example values.

Create a port ACL rule

1. Get an existing port ACL.

```
PS C:\> $portACL = Get-SCPortACL -Name "RDPAccess"
```

2. Create a port ACL rule.

```
PS C:\> New-SCPortACLRule -Name "AllowRDPAccess" -PortACL $portACL -Description "Allow RDP Rule from a subnet" -Action Allow -Type Inbound -Priority 110 -Protocol Tcp -LocalPortRange 3389 - RemoteAddressPrefix 10.184.20.0/24
```

NOTE

- Priority range for SDN port ACL rules: 1 – 64500.
- Only TCP/UDP/Any protocol parameters are supported for creating ACL rules.

Attach an ACL to a virtual network adapter

1. Get the virtual network adapter.

```
PS C:\> $vm = Get-SCVirtualMachine -Name "TenantVM"  
PS C:\> $adapter = Get-SCVirtualNetworkAdapter -VM $vm
```

2. Attach an existing port ACL to the virtual network adapter.

```
PS C:\> $portACL = Get-SCPortACL -Name "RDPAccess"  
PS C:\> Set-SCVirtualNetworkAdapter -VirtualNetworkAdapter $adapter -PortACL $portACL
```

NOTE

You can also attach a port ACL while creating the virtual network adapter through **New-SCVirtualNetworkAdapter** cmdlet. [Learn more](#).

Detach a port ACL from a virtual network adapter

1. Get the virtual network adapter that you want to detach the port ACL from.

```
PS C:\> $vm = Get-SCVirtualMachine -Name "TenantVM"  
PS C:\> $adapter = Get-SCVirtualNetworkAdapter -VM $vm
```

2. Detach the port ACL from the virtual network adapter.

```
PS C:\> Set-SCVirtualNetworkAdapter -VirtualNetworkAdapter $adapter -RemovePortACL
```

Attach an ACL to a VM subnet

1. Get the VM subnet to attach the ACL.

```
PS C:\> $vmSubnet = Get-SCVMSubnet -Name "Tenant Subnet"
```

2. Attach an existing port ACL to the VM subnet.

```
PS C:\> Set-SCVMSubnet -VMSubnet $vmSubnet -PortACL $portACL
```

NOTE

You can also attach a port ACL while creating VM subnet through **New-SCVMSubnet cmdlet**. [Learn more](#).

Detach a port ACL from a VM subnet

1. Get the VM subnet that you want to detach the port ACL from.

```
PS C:\> $vmSubnet = Get-SCVMSubnet -Name "Tenant Subnet"
```

2. Detach the port ACL from the VM subnet.

```
PS C:\> Set-SCVMSubnet -VMSubnet $vmSubnet -RemovePortACL
```

Remove a port ACL rule

1. Get the port ACL rule to remove.

```
PS C:\> $portACLRule = Get-SCPortACLRule -Name "AllowRDPAccess"
```

2. Remove the port ACL rule.

```
PS C:\> Remove-SCPortACLRule -PortACLRule $portACLRule
```

Remove a port ACL

1. Get the Port ACL that you want to remove.

```
PS C:\> $portACL = Get-SCPortACL -Name "RDPAccess"
```

2. Remove the port ACL.

```
PS C:\> Remove-SCPortACL -PortACL $portACL
```


Control SDN virtual network bandwidth with QoS

5/30/2019 • 3 minutes to read

This article provides information about how to configure the Quality of Service (QoS) settings for SDN virtual networks, in System Center - Virtual Machine Manager (VMM). Through this configuration, you can limit the bandwidth of the traffic flowing in or out of a virtual network adapter (vNIC) by specifying the minimum reserved bandwidth or maximum bandwidth for the adapter.

- Service providers can prevent a high-traffic VM from blocking other VMs.
- Tenants can get the minimum reserved bandwidth regardless of the network traffic.

For detailed information about the SDN QoS settings available in Network Controller, see [Configure Quality of Service \(QoS\) for a Tenant VM Network Adapter](#).

QoS settings are managed in VMM through virtual network adapter port profiles associated with a port classification. Port profile settings enable uniformity and ease of application across multiple adapters.

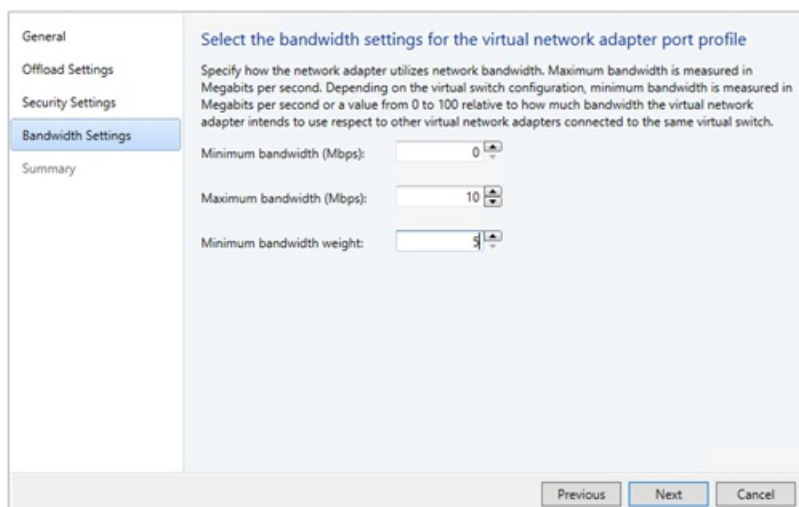
Before you start

Ensure you have the network controller deployed in the SDN fabric. [Learn more](#).

Create virtual network adapter port profiles

To create virtual network adapter port profile, use the following procedure:

1. Go to **Fabric > Port Profiles**.
2. Right-click **Port Profile** and select **Create Hyper -V Port Profile**. Provide a name and optional description. For example, name can be **TenantA**.
3. In **General**, select **Virtual Network Adapter Port Profile**.
4. Go to **Bandwidth Settings**.



The screenshot shows the 'Bandwidth Settings' tab of a configuration window. The left sidebar has tabs for General, Offload Settings, Security Settings, Bandwidth Settings (selected), and Summary. The main area is titled 'Select the bandwidth settings for the virtual network adapter port profile'. Below the title is a descriptive paragraph: 'Specify how the network adapter utilizes network bandwidth. Maximum bandwidth is measured in Megabits per second. Depending on the virtual switch configuration, minimum bandwidth is measured in Megabits per second or a value from 0 to 100 relative to how much bandwidth the virtual network adapter intends to use respect to other virtual network adapters connected to the same virtual switch.' There are three input fields: 'Minimum bandwidth (Mbps)' with a value of 0, 'Maximum bandwidth (Mbps)' with a value of 10, and 'Minimum bandwidth weight' with a value of 5. At the bottom right are 'Previous', 'Next', and 'Cancel' buttons.

5. In **Bandwidth Settings**, enter the minimum and maximum bandwidth values that you want to apply to the selected vNIC. Maximum bandwidth (Mbps) and Minimum bandwidth (Weight) are used to reserve a portion of the outbound bandwidth for the vNICs.

- **Maximum bandwidth (Mbps):** Specify the maximum outbound bandwidth that can be used by this

vNIC. This is irrespective of the bandwidth mode configured on the switch. A value of 0 implies that the maximum is not configured.

- **Minimum bandwidth (Weight):** Specify a weighted portion of the outbound bandwidth that you want to reserve for this vNIC. This will reserve a weighted portion of the total physical NIC(s) bandwidth based on the total weights reserved by all other vNICs for a particular switch.

6. In **Summary**, view the settings and click **Finish**.

NOTE

- Minimum bandwidth (Mbps) is not supported by VMM for SDN.
- Limiting the inbound bandwidth of the vNIC is not supported by VMM.
- Minimum bandwidth (Weight) sets the **outboundReservedValue** setting on the network controller and maximum bandwidth (Mbps) sets the **outboundMaximumMbps** setting.
- Ensure that the minimum and maximum values for bandwidth are within the bandwidth range of the physical NIC(s) that the logical switch is deployed on. Otherwise the bandwidth allocation request will be rejected.

Create port classification

To create a port classification:

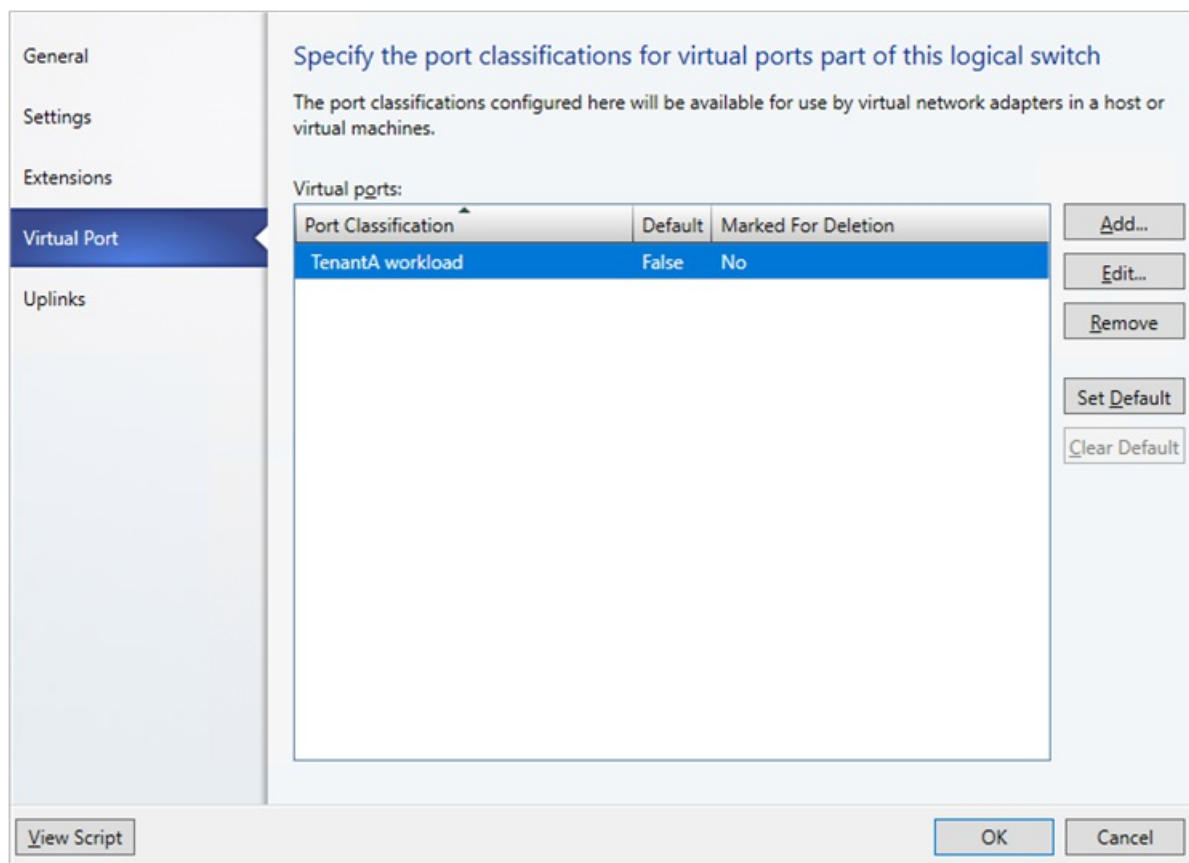
1. Go to **Fabric > Networking**.
2. Right-click **Port Classifications > Create Port Classification**.
3. In the **Create Port Classification** wizard, in **General**, give a name to the classification and click **Finish**.

As an example, we created a port classification with the name **TenantA workload**, which we use later.

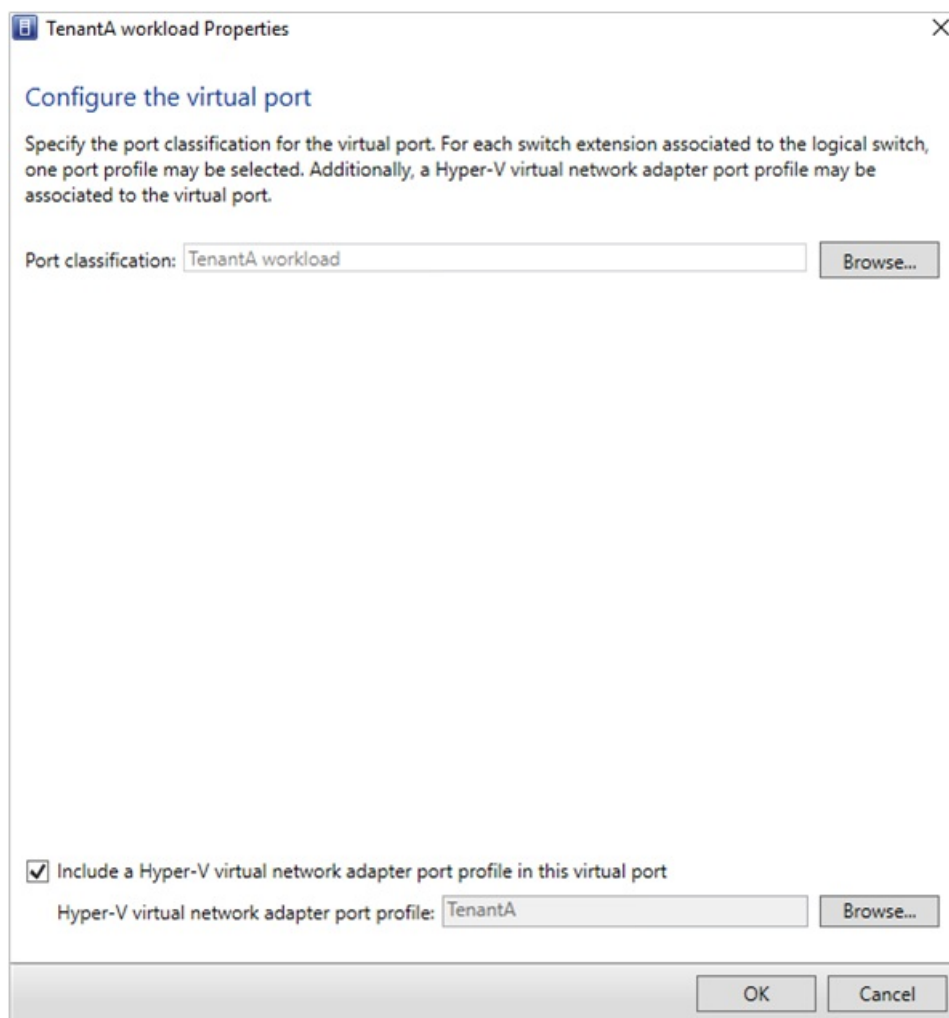
Associate port classifications to logical switch

To associate a port classification to the logical switch, use the following procedure:

1. Go to **Fabric > Networking > Logical Switches**.
2. Right-click the selected logical switch to view its **Properties**.
3. Select Virtual Ports , click **Add** to open the **Add Virtual Port** dialog.



4. **Browse** to select the port classification (as an example: **TenantA Workload**) and click **OK**.
5. Select the **Include the Hyper-V virtual network adapter port profile** option and click **Browse**. Select the vNIC port profile, (as an example, **TenantA** that you created in the previous procedure) click **OK**.

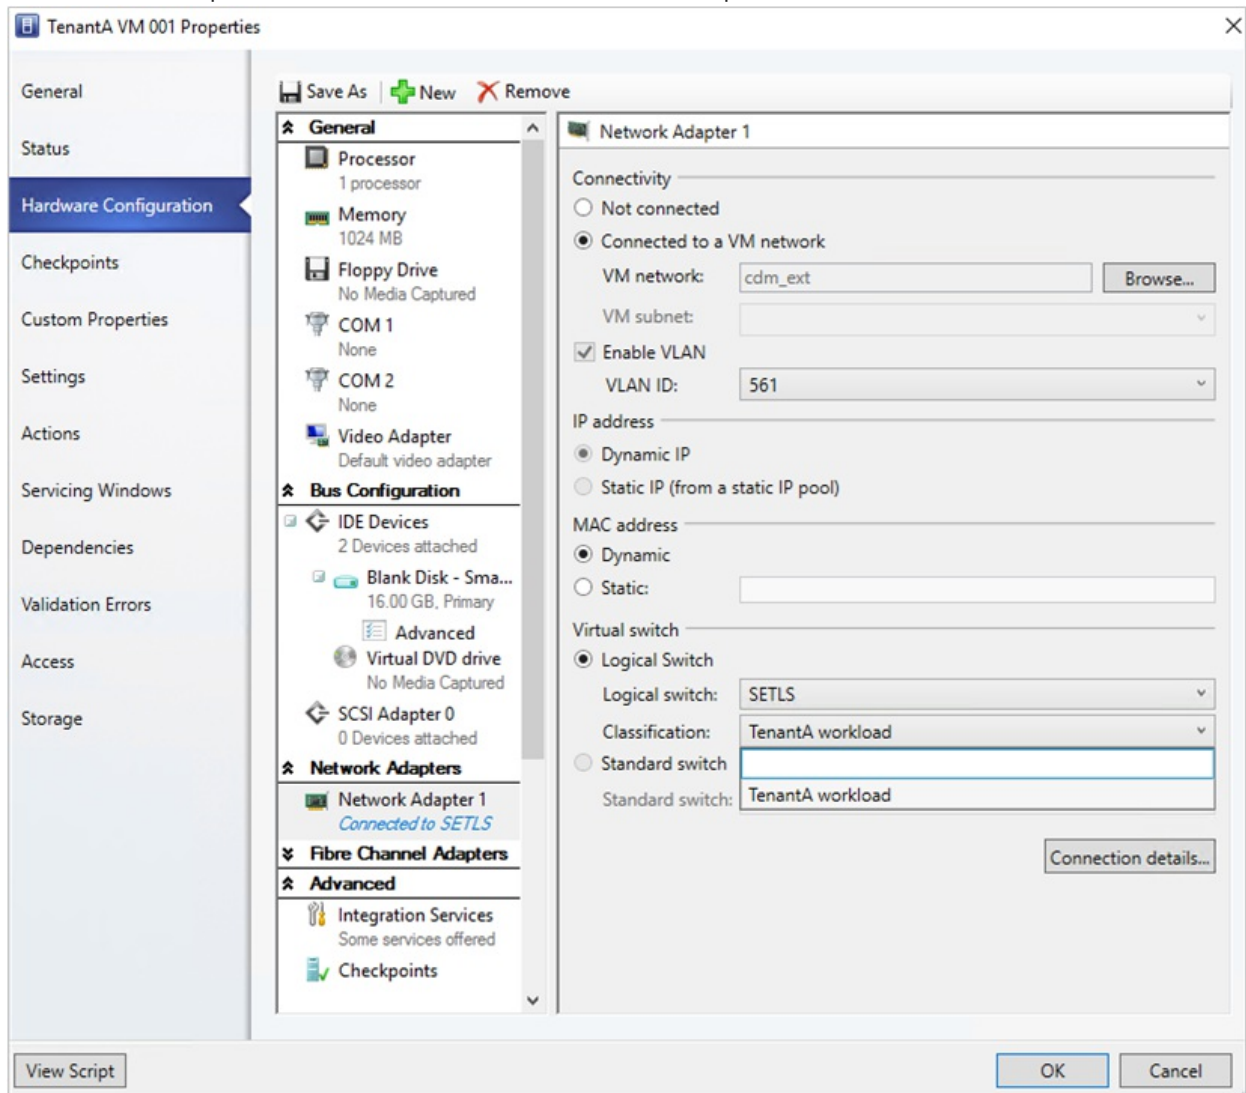


6. Repeat the above steps to add the port classifications required for this logical switch.

Apply port classification to virtual network adapter

To apply a port classification (and hence the QoS settings) to a virtual network adapter, use the following procedure:

1. Go to **VMs and Services**, right-click the selected VM and open its **Properties**.
2. Go to the **Hardware Configuration** page and select the virtual network adapter.
3. Select the desired port classification from the **Classification** drop-down list.



NOTE

- Only the classifications available for the logical switch are displayed here. You can also apply a port classification while creating a VM.
- You can also apply a port classification to the host virtual network adapter via the host properties Virtual Switches page.

Load balance network traffic in an SDN fabric using VMM

5/6/2019 • 5 minutes to read

In Software Define Networking (SDN) fabric, you can use the Software Load Balancer (SLB) to evenly distribute the network traffic among the workloads managed by the Cloud Service Providers (CSPs) and the tenants. Learn more about [Software Load Balancer \(SLB\)](#).

System Center Virtual Machine Manager (VMM) supports the following two scenarios of load balancing:

- **Load balancing the external network traffic:** This scenario includes load balancing the network traffic coming from an external network. Load balancing is done through a public Virtual IP (VIP). The workload VMs can be connected to a **no isolation** VM network or to a **network virtualization** enabled VM network.
- **Load balancing the internal network traffic** (Internal load balancing): This scenario includes load balancing the network traffic of workloads in the VM network that is enabled for **network virtualization**. Load balancing is done through a VIP.

NOTE

VMM 2016 supports configuration of load balancing by using PowerShell. Configuration through Service templates is currently not supported.

This article provides information about how to configure the load balancing of workload VMs by using PowerShell.

Before you start

Ensure the following:

- [SDN network controller](#) and the SDN [software load balancer](#) are deployed.
- Required VMs are created on appropriate networks and are identified for load balancing.

Create a VIP template

Use the following procedures to configure the VIPs to load balance the workload VMs (DIPs).

1. In VMM console, click **Fabric > VIP Templates**. Right-click and select **Create VIP Template**.

Name	<h3>Specify a name and description for this template</h3> <p>A virtual IP template contains load balancer-related configuration settings for a specific type of network traffic. The virtual port is the port that is used for the type of network traffic that you want to load balance.</p> <p>Template name: <input type="text" value="SLB80"/></p> <p>Description: <input type="text" value="Load balance network traffic to the website"/></p> <p>Virtual IP port: <input type="text" value="80"/></p> <p>Backend port: <input type="text" value="80"/></p>
Type	
Protocol	
Persistence	
Load Balancing	
Health Monitors	
Summary	

- In the **Name** page, provide the name and optional description. Specify the **Virtual IP Port** that will be used as front end port for the network traffic. In the **Backend port**, specify the port on which the backend workloads are hosted. Click **Next**.

Name	<h3>Specify a template type</h3> <p>Select the type of template to create that best matches your environment.</p> <p><input type="radio"/> Generic A generic template can be used on any supported load balancer.</p> <p><input checked="" type="radio"/> Specific A virtual IP template can be used only on a load balancer of a specific manufacturer and model.</p> <p>Manufacturer: <input type="text" value="Microsoft"/></p> <p>Model: <input type="text" value="Microsoft Network Controller"/></p>
Type	
Protocol	
Persistence	
Load Balancing	
Health Monitors	
Summary	

- In **Specify a Template type**, select Specific. In **Manufacturer**, select Microsoft. In **Model**, select Microsoft Network Controller. Click **Next**.

Name
Type
Protocol
Persistence
Load Balancing
Health Monitors
Summary

Specify protocol options

Select the protocol for which you want to create the virtual IP template. If you select HTTPS, you must specify whether to terminate the HTTPS traffic at the load balancer or to pass it through to the virtual machine.

☐ HTTP
☐ **HTTPS passthrough**
Passes through HTTPS traffic without decrypting at load balancers.
☐ **HTTPS terminate**
Terminates HTTPS traffic at load balancers.
Certificate subject name:
☒ Re-Encrypt
Re-encrypt HTTPS traffic from the load balancer to the server.
☒ **Custom**
Protocol name:

PreviousNextCancel

4. In **Specify Protocol options**, select the Custom option and type TCP/UDP in Protocol name. Click **Next**.

NOTE

SLB supports only layer 4 load balancing of TCP/UDP traffic. HTTP/ HTTPS protocols are not supported.

5. For **Persistence** and **Load Balancing**, use the defaults. Click **Next**.

Name
Type
Protocol
Persistence
Load Balancing
Health Monitors
Summary

Specify a health monitor

Specify one or more requests that you want to send at regular intervals to verify that a load balancer is available. Adding a health monitor is optional.

Health monitors:

	Protocol	Request	Response	Interval	Time-Out	Retries	
	Http	/	200	5	2	3	Insert Delete

PreviousNextCancel

NOTE

SLB doesn't support persistence and different Load balancing methods. By default, SLB uses Round Robin algorithm for load balancing.

6. In **Health Monitors**, optionally, insert appropriate values. Click **Next**.

Options: TCP and HTTP are supported as health monitor protocols if SLB is used.

7. Verify the settings, and click **Finish**.

Once the template is created, you can find this under **Fabric > VIP Templates**.

Create a VIP using PowerShell

Create a VIP by using the following example scripts. Select the script based on the type of network traffic that you want to load balance:

- [Script for creating a VIP to load balance external network traffic](#)
- [Script for creating a VIP to load balance internal network traffic](#)

Script for creating VIP to load balance external network traffic

The following example script creates the VIP from a public IP network to load balance the workload VMs WGB-001 & WGB-002 on port 80.

The workload VMs can be connected to a **no isolation** network or **network virtualization** enabled VM networks.

NOTE

- In the script parameters section, substitute the actual values that match your test environment for the samples used in this script.
- Ensure you run the script on a VMM server, or on a computer running the VMM console.

```
param(

[Parameter(Mandatory=$false)]
# Name of the Network Controller Network Service
# This value should be the name you gave the Network Controller service
# when you on-boarded the Network Controller to VMM
$LBServiceName = "NC",

[Parameter(Mandatory=$false)]
# Name of the workload VMs you want to load balance.
$VipMemberVMNames = @( "WGB-001", "WGB-002"),

[Parameter(Mandatory=$false)]
# Name of the VIP VM Network
$VipNetworkName = "PublicVIP",

[Parameter(Mandatory=$false)]
# VIP address you want to assign from the VIP VM Network IP pool.
# Pick any VIP that falls within your VIP IP Pool range.
$VipAddress = "44.15.10.253",

[Parameter(Mandatory=$false)]
# The name of the VIP template you created via the VMM Console.
$VipTemplateName = "WebsiteHTTP",

[Parameter(Mandatory=$false)]
# Arbitrary but good to match the VIP you're using.
$VipName = "scvmm_44_15_10_253_80"

)
```

Import-Module virtualmachinemanagement


```

import-module virtualmachinemanagement

$lb = Get-SCLoadBalancer | where { $_.Service.Name -like $LBServiceName};
$vipNetwork = get-scvmnetwork -Name $VipNetworkName;

$vipMemberNics = @();
foreach ($vmName in $VipMemberVMNames)
{
    $vm = get-scvirtualmachine -Name $vmName;
    # if ($vm.VirtualNetworkAdapters[0].VMNetwork.ID -ne $vipNetwork.ID)
    # {
    #     $vm.VirtualNetworkAdapters[0] | set-scvirtualnetworkadapter -VMNetwork $vipNetwork;
    # }

    $vipMemberNics += $vm.VirtualNetworkAdapters[0];
}

$existingVip = get-scloadbalancervip -Name $VipName
if ($existingVip -ne $null)
{
    # foreach ($mem in $existingVip.VipMembers)
    # {
    #     $mem | remove-scloadbalancervipmember;
    # }

    $existingVip | remove-scloadbalancervip;
}

$vipt = get-scloadbalancerviptemplate -Name $VipTemplateName;

$vip = New-SCLoadBalancerVIP -Name $VipName -LoadBalancer $lb
-IPAddress $VipAddress -LoadBalancerViptemplate $vipt
-FrontEndVMNetwork $vipNetwork
-BackEndVirtualNetworkAdapters $vipMemberNics;
Write-Output "Created VIP " $vip;

$vip = get-scloadbalancervip -Name $VipName;
Write-Output "VIP created successfully " $vip;

```

Script for creating VIP to load balance internal network traffic

For the following example script, we created a new VIP template by name ILB-VIP-Template for load balancing the port 443 using the procedure detailed in the [previous section](#). The script creates a VIP from tenant VM network to load balance the VMs ILB-001 & ILB-002, which are part of the same tenant VM network.

NOTE

- In the Internal load balancing scenario, the VIP comes from the tenant VM network. So, the **\$VipNetworkName** is the same tenant VM network name where the VMs are connected. Ensure that the tenant VM network IP pool has the IPs reserved for VIPs. The **\$VipAddress** is one of the IP addresses from the reserved VIPs.
- In the script parameters section, substitute the actual values that match your test environment for the samples used in this script.
- Ensure you run the script on a VMM server, or on a computer running the VMM console.

```

param(

[Parameter(Mandatory=$false)]
# Name of the Network Controller Network Service
# This value should be the name you gave the Network Controller service
# when you on-boarded the Network Controller to VMM
$LBServiceName = "NC",

[Parameter(Mandatory=$false)]
# Name of the workload VMs you want to load balance.

```

```

$VipMemberVMNames = @("ILB-001","ILB-002"),

[Parameter(Mandatory=$false)]
# Name of the VIP VM Network
$VipNetworkName = "TenantNetwork",

[Parameter(Mandatory=$false)]
# VIP address you want to assign from the VIP VM Network IP pool.
# Pick any VIP that falls within your VIP IP Pool range.
$VipAddress = "192.168.100.75",

[Parameter(Mandatory=$false)]
# The name of the VIP template you created via the VMM Console.
$VipTemplateName = "ILB-VIP-Template",

[Parameter(Mandatory=$false)]
# Arbitrary but good to match the VIP you're using.
$VipName = "scvmm_192_168_100_75_443"

)

Import-Module virtualmachinemanager

$lb = Get-SCLoadBalancer | where { $_.Service.Name -like $LBSERVICEName};
$vipNetwork = get-scvmmnetwork -Name $VipNetworkName;

$vipMemberNics = @();
foreach ($vmName in $VipMemberVMNames)
{
    $vm = get-scvirtualmachine -Name $vmName;
    # if ($vm.VirtualNetworkAdapters[0].VMNetwork.ID -ne $vipNetwork.ID)
    # {
    #     $vm.VirtualNetworkAdapters[0] | set-scvirtualnetworkadapter -VMNetwork $vipNetwork;
    # }

    $vipMemberNics += $vm.VirtualNetworkAdapters[0];
}

$existingVip = get-scloadbalancervip -Name $VipName
if ($existingVip -ne $null)
{
    # foreach ($mem in $existingVip.VipMembers)
    # {
    #     $mem | remove-scloadbalancervipmember;
    # }

    $existingVip | remove-scloadbalancervip;
}

$vipt = get-scloadbalancerviptemplate -Name $VipTemplateName;

$vip = New-SCLoadBalancerVIP -Name $VipName -LoadBalancer $lb
-IPAddress $VipAddress -LoadBalancerVIPTemplate $vipt
-FrontEndVMNetwork $vipNetwork
-BackEndVirtualNetworkAdapters $vipMemberNics;
Write-Output "Created VIP " $vip;

$vip = get-scloadbalancervip -Name $VipName;
Write-Output "VIP created successfully " $vip;

```

Set up NAT for traffic forwarding in the SDN infrastructure

3/14/2019 • 2 minutes to read

This article describes how to set up Network Address Translation (NAT) for traffic forwarding in a software-defined network (SDN) infrastructure set up in the System Center Virtual Machine Manager (VMM) fabric.

NAT allows virtual machines (VMs) in an isolated SDN virtual network to obtain external connectivity. VMM configures a Virtual IP (VIP) to forward the traffic to and from an external network.

The following two NAT types are supported by VMM.

- **Outbound NAT** - Forwards the VM network traffic from a virtual network to external destinations.
- **Inbound NAT** - Forwards the external traffic to a specific VM in a virtual network.

This article provides information about how to configure a NAT connection for SDN virtual networks using VMM.

Before you start

Ensure the following:

- [SDN network controller](#) and the SDN [software load balancer](#) are deployed.
- An SDN VM network with network virtualization is created.

Create a NAT connection

Use the following procedure:

1. In VMM console, click **VMs and Services > VM Networks**. Right-click the selected VM network for which you want to create the NAT connection, and select **Properties**.
2. Click **Connectivity** on the wizard page displayed.

TenantNetwork Properties

Name

VM Subnets

Connectivity

Network Address...

Access

Connectivity

By default, network virtualization provides an isolated network. Select the options below to add external connectivity to this network.

☐ Connect to another network through a VPN tunnel

☐ Enable Border Gateway Protocol (BGP)

☒ Connect directly to an additional logical network

☐ Direct routing

☒ Network address translation (NAT)

Select the gateway device to use for this connection.

Gateway Device: NC

Device capabilities: Site to Site VPN
Border Gateway Protocol (BGP)
Direct routing

Remaining capacity: 50 VM networks
250 VPN tunnels

View Script

OK Cancel

3. In **Connectivity**, select **Connect directly to an additional logical network** and select **Network address translation (NAT)** under this option.

TenantNetwork Properties

Name

VM Subnets

Connectivity

Network Address...

Access

Network address translation (NAT)

In NAT all traffic to external network is routed through single IP, select the IP address pool and IP below. If you want VMM to automatically pick any available IP from the selected address pool, leave the IP address column empty.

IP address pool: PublicVIP (44.55.10.4 - 44.55.10.254)

IP address: 44.55.10.25

Specify network address translation (NAT) rules:

Name	Protocol	Incoming Port	Destination IP	Destination P...

Add

Remove

View Script

OK Cancel

4. In the **IP address pool**, choose the IP pool from which the VIP should come from. In IP address, choose an IP address from the pool selected. Click **OK**.

A NAT connection will be created for this VM network.

NOTE

- Along with the NAT connection, this procedure also creates a default Outbound NAT rule that enables the outbound connectivity for the VM network.
- To enable inbound connectivity and forward an external traffic to a specific VM, you must add NAT rules to the NAT connection.

Add rules to a NAT connection

Use the following procedure:

1. In VMM console, click **VMs and Services > VM Networks**. Right-click the selected VM network and select **Properties**.
2. Click **Network Address Translation** on the wizard.

The screenshot shows the 'TenantNetwork Properties' dialog box with the 'Network Address Translation (NAT)' tab selected. The left sidebar contains a tree view with 'Name', 'VM Subnets', 'Connectivity', 'Network Address...', and 'Access'. The main area displays the NAT configuration. It includes a text box for 'IP address pool' set to 'PublicVIP (44.55.10.4 - 44.55.10.254)' and an 'IP address' field set to '44.55.10.25'. Below this is a table titled 'Specify network address translation (NAT) rules:'. The table has five columns: 'Name', 'Protocol', 'Incoming Port', 'Destination IP', and 'Destination P...'. One rule is listed: 'RDP' with 'TCP' protocol, '3389' incoming port, '192.168.10.23' destination IP, and '3389' destination port. To the right of the table are 'Add' and 'Remove' buttons. At the bottom of the dialog are 'View Script', 'OK', and 'Cancel' buttons.

TenantNetwork Properties

Network address translation (NAT)

In NAT all traffic to external network is routed through single IP, select the IP address pool and IP below. If you want VMM to automatically pick any available IP from the selected address pool, leave the IP address column empty.

IP address pool: PublicVIP (44.55.10.4 - 44.55.10.254)

IP address: 44.55.10.25

Specify network address translation (NAT) rules:

Name	Protocol	Incoming Port	Destination IP	Destination P...
RDP	TCP	3389	192.168.10.23	3389

Buttons: Add, Remove, View Script, OK, Cancel

3. Under **Specify network address translation (NAT) rules**, click **Add**. Type the following details as appropriate:
 - **Name** – Name for the inbound NAT rule.
 - **Protocol** – Inbound network traffic protocol. TCP/UDP are supported.
 - **Incoming Port** – Port number that you want to use along with the VIP to access the VM.
 - **Destination IP** – IP address of the VM to which you want to direct the external traffic.
 - **Destination Port** – Port number on the VM, the external traffic should be forwarded to.
4. Click **OK**.

NOTE

Multiple NAT rules can be created to forward the traffic to multiple VMs that are part of the VM network.

Remove a NAT rule

Use the following procedure:

1. In VMM console, click **VMs and Services > VM Networks**. Right-click the selected VM network and select **Properties**.
2. Click **Network Address Translation** on the wizard.
3. Select the NAT rule that you want to remove, Click **Remove** and then click **OK**.

Remove a NAT connection

1. In VMM console, click **VMs and Services > VM Networks**. Right-click the selected VM network and select **Properties**.
2. Click **Connectivity** on the wizard.
3. Clear the option **Connect directly to an additional logical network** and click **OK**.

Route traffic across networks in the SDN infrastructure

5/30/2019 • 14 minutes to read

This article describes how to route traffic across networks in a software-defined network (SDN) infrastructure set up in the System Center Virtual Machine Manager (VMM) fabric.

An SDN RAS gateway enables you to route network traffic between physical and virtual networks, regardless of where the resources are located. SDN RAS gateway is multitenant, Border Gateway Protocol (BGP) capable and supports connectivity using Site-to-Site virtual private network (VPN) using IPsec or Generic Routing Encapsulation (GRE) or Layer 3 Forwarding. [Learn more.](#)

Before you start

Ensure the following:

- SDN [Network Controller](#), SDN [Software Load Balancer](#), and [SDN RAS gateway](#) are deployed.
- An SDN VM network with network virtualization is created.

Configure Site-to-Site VPN connections using VMM

A site-to-site VPN connection allows you to securely connect two networks at different physical locations by using Internet.

For Cloud Service Providers (CSPs) that host many tenants in their datacenter, SDN RAS gateway provides a multi-tenant gateway solution that allows your tenants to access and manage their resources over Site-to-Site VPN connections from remote sites, which in turn allows network traffic between virtual resources in your datacenter and their physical network.

Configure IPSec connection

Use the following procedure:

1. Select the VM Network that you want to configure a Site-to-Site IPSec connection, and click **Connectivity**.
2. Select **Connect to another network through a VPN tunnel**. Optionally, to enable BGP peering in your datacenter, select **Enable Border Gateway Protocol (BGP)**.
3. Select the network controller service for the gateway device.
4. Select the **VPN Connections > Add > Add IPSec Tunnel**.
5. Type a subnet as shown in the following diagram. This subnet is used to route packets out of the VM Network. You do not need to pre-configure this subnet in your datacenter.

tenant1 Properties

Name

VM Subnets

Connectivity

VPN Connections

Access

Routing Subnet

A subnet is required for routing packets out of the VM network.

Subnet: 10.252.254.0/29

Specify VPN connections:

+ Add - Remove

TEST_IPSEC 10.184.109.200	IPsec tunnel
Authentication Run As account	Name: TEST_IPSEC
Routes 0 routes	VPN endpoint
Advanced PSKOnly	Remote Endpoint: 10.184.109.200
	Bandwidth
	<input type="checkbox"/> Limit bandwidth available for this VPN connection
	Maximum Incoming (Mbps): 0
	Maximum Outgoing (Mbps): 0

View Script

OK Cancel

- Type a name for the connection, and the IP address of the remote endpoint. Optionally, configure the bandwidth.
- In **Authentication**, select the type of authentication you want to use. If you choose to authenticate by using a Run as account, create a user account with a user name, and the IPsec key as the password for the account.
- In **Routes**, type all the remote subnets that you want to connect to. If you have selected **Enable Border Gateway Protocol (BGP)** in the **Connectivity** page, routes are not required.
- On the **Advanced** tab, accept the default settings.
- If you have selected **Enable Border Gateway Protocol (BGP)** in the Connectivity page, then you can fill out your ASN, peer BGP IP, and its ASN on the **Border Gateway Protocol** wizard page as shown below:

tenant2 Properties

Name

VM Subnets

Connectivity

VPN Connections

Border Gateway Protocol...

Access

Border Gateway Protocol (BGP)

BGP is a protocol that enables automatic exchanges of routes between individual networks. Enable BGP only if your environment is configured for it.

Autonomous system number (ASN): 134

BGP peers:

IP Address	Autonomous System Number	
10.200.100.55	135	

Add

Remove

View Script

OK Cancel

11. To validate the connection, try to ping the remote endpoint IP address from one of the virtual machines on your VM network.

Configure GRE tunneling

GRE tunnels enable connectivity between tenant virtual networks and external networks. Since the GRE protocol is lightweight and support for GRE is available on most of the network devices, it becomes an ideal choice for tunneling where encryption of data is not required. GRE support in Sit-to-Site (S2S) tunnels facilitates traffic forwarding between tenant virtual networks and tenant external networks.

Use the following procedure:

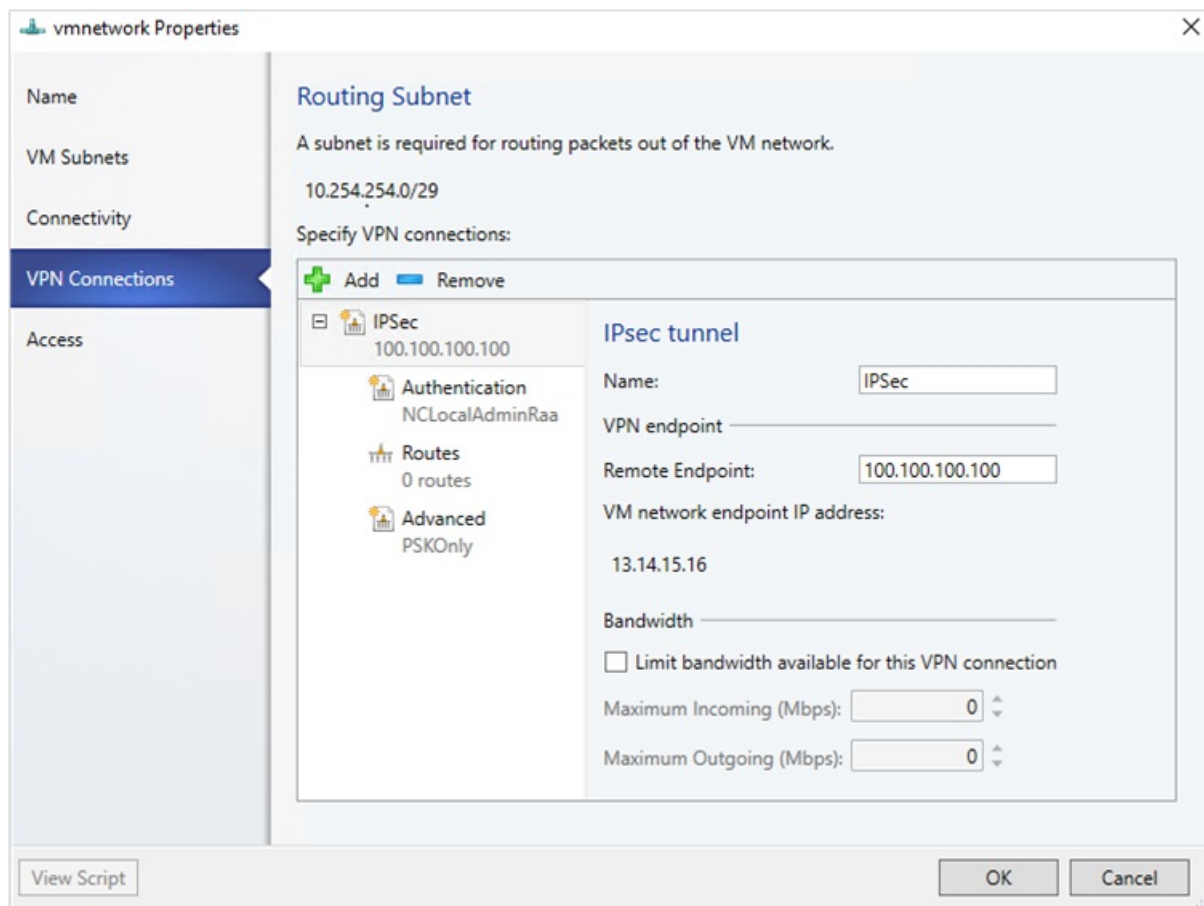
1. Select the VM network where you want to configure a S2S GRE connection, and click **Connectivity**.
2. Select **Connect to another network through a VPN tunnel**. Optionally, to enable BGP peering in your datacenter, select **Enable Border Gateway Protocol (BGP)**.
3. Select the Network Controller Service for the Gateway Device.
4. Select **VPN Connections > Add > Add GRE Tunnel**.
5. Type a subnet as shown in the following diagram. This subnet is used to route packets out of the VM network. This subnet doesn't need to be preconfigured in your datacenter.

The screenshot shows the 'tenant1 Properties' dialog box with the 'VPN Connections' tab selected. The 'Routing Subnet' section displays a subnet of 10.252.254.0/29. Below this, the 'Specify VPN connections' section contains a table with one entry: 'TEST_GRE' with IP '10.184.109.200'. To the right of the table, the 'GRE tunnel' configuration fields are visible, including Name (TEST_GRE), Remote Endpoint (10.184.109.200), GRE Key (1234), Internal IPv4 Address, Internal IPv6 Address, Bandwidth, and a checkbox for 'Limit bandwidth available for this VPN connection'.

6. Type a connection name, and specify the IP address of the remote endpoint.
7. Type the **GRE key**.
8. Optionally, you can complete the other fields on this screen, these values aren't needed to set up a connection.
9. In **Routes**, add all the remote subnets that you want to connect to. If you selected **Enable Border Gateway Protocol (BGP)** in **Connectivity**, you can leave this screen blank and instead complete your ASN, peer BGP IP, and ASN fields on the **Border Gateway Protocol** tab.
10. You can use the defaults for the remaining settings.
11. To validate the connection, try to ping the remote endpoint IP address from one of the virtual machines on the VM network.

Configure IPsec and GRE connections on the remote site

On the remote peer device, use the **VM network endpoint IP address** from the VMM UI as destination Address while setting up the IPsec\GRE connection.



Configure L3 forwarding

L3 forwarding enables connectivity between the physical infrastructure in the datacenter and the virtualized infrastructure in the Hyper-V network virtualization cloud.

Using L3 forwarding, tenant network virtual machines can connect to a physical network through the Windows Server 2016 SDN Gateway, which is already configured in an SDN environment. In this case, the SDN gateway acts as a router between the virtualized network and the physical network.

To learn more, check these articles: [Windows server gateway as a forwarding gateway](#) and [RAS gateway high availability](#).

Ensure the following before you attempt to configure L3:

- Ensure you're logged on as an administrator on the VMM server.
- You must configure a unique next-hop logical network, with unique VLAN ID, for each Tenant VM network for which L3 forwarding needs to be set up. There must be 1:1 mapping between a tenant network and corresponding physical network (with unique VLAN ID).

Use the following steps to create the next-hop logical network in SCVMM:

1. On the VMM console, select **Logical Networks**, right-click, and select **Create Logical Network**.
2. In the **Settings** page, choose **One connected network** and select the checkbox for **Create a VM network with the same name to allow virtual machines to access this logical network directly** and **Managed by Microsoft Network Controller**

3. Create an IP Pool for this new logical network.

IP address from this pool is required in the script for setting up L3 forwarding.

The following table provides examples of dynamic and static L3 connections.

PARAMETER	DETAILS/EXAMPLE VALUES
L3VPNConnectionName	User-defined name for the L3 forwarding network connection. Example: Contoso_L3_GW
VmNetworkName	Name of the tenant virtual network that's reachable over L3 network connection. This network must exist when running the script. Example: ContosoVMNetwork
NextHopVMNetworkName	User-defined name for the next hop VM network, which was created as a prerequisite. This represents the physical network that wants to communicate with the tenant VM network. This network must exist when running this script. Example: Contoso_L3_Network
LocalIPAddresses	IP addresses to be configured on the SDN gateway L3 network interface. This IP address must belong to the next hop logical network you created. You must also provide the subnet mask. Example: 10.127.134.55/25
PeerIPAddresses	IP address of the physical network gateway, reachable over L3 logical network. This IP address must belong to the next hop logical network you created in the prerequisites. This IP will serve as the next hop once traffic destined to the physical network from the tenant VM network reaches the SDN gateway. Example: 10.127.134.65
GatewaySubnet	Subnet to be used for routing between HNV gateway and tenant virtual network. You can use any subnet, ensure that it does not overlap with the next hop logical network. Example: 192.168.2.0/24
RoutingSubnets	Static routes that need to be on the L3 interface of the HNV gateway. These routes are for the physical network subnets, which should be reachable from the tenant VM network over the L3 connection.
EnableBGP	Option to enable BGP. Default: false.
TenantASNRoutingSubnets	ASN number of the tenant gateway, only if BGP is enabled.

Run the following script to set up L3 forwarding. Refer to the table above to check what each script parameter identifies.

```
``powershell
param (
    [Parameter(Mandatory=$true)]
    # Name of the L3 VPN connection
    $L3VPNConnectionName,
    [Parameter(Mandatory=$true)]
    # Name of the VM network to create gateway
    $VmNetworkName,
    [Parameter(Mandatory=$true)]
    # Name of the Next Hop one connected VM network
```

```

# used for forwarding
$NextHopVmNetworkName,
[Parameter(Mandatory=$true)]
# IPAddresses on the local side that will be used
# for forwarding
# Format should be @"10.10.10.100/24"
$LocalIPAddresses,
[Parameter(Mandatory=$true)]
# IPAddresses on the remote side that will be used
# for forwarding
# Format should be @"10.10.10.200"
$PeerIPAddresses,
[Parameter(Mandatory=$false)]
# Subnet for the L3 gateway
# default value 10.254.254.0/29
$GatewaySubnet = "10.254.254.0/29",
[Parameter(Mandatory=$false)]
# List of subnets for remote tenants to add routes for static routing
# Format should be @"14.1.20.0/24","14.1.20.0/24";
$RoutingSubnets = @(),
[Parameter(Mandatory=$false)]
# Enable BGP in the tenant space
$EnableBGP = $false,
[Parameter(Mandatory=$false)]
# ASN number for the tenant gateway
# Only applicable when EnableBGP is true
$TenantASN = "0"
)

# Import SC-VMM PowerShell module
Import-Module virtualmachinemanager

# Retrieve Tenant VNET info and exit if VM Network not available
$VmNetwork = Get-SCVMNetwork -Name $VmNetworkName;
if ($VmNetwork -eq $null)
{
    Write-Verbose "VM Network $VmNetworkName not found, quitting"
    return
}

# Retrieve L3 Network info and exit if VM Network not available
$nextHopVmNetwork = Get-SCVMNetwork -Name $NextHopVmNetworkName;
if ($nextHopVmNetwork -eq $null)
{
    Write-Verbose "Next Hop L3 VM Network $NextHopVmNetworkName not found, quitting"
    return
}

# Retrieve gateway Service and exit if not available
$gatewayDevice = Get-SCNetworkGateway | Where {$_ .Model -Match "Microsoft Network Controller"};
if ($gatewayDevice -eq $null)
{
    Write-Verbose "Gateway Service not found, quitting"
    return
}

# Retrieve Tenant Virtual Gateway info
$VmNetworkGatewayName = $VmNetwork.Name + "_Gateway";
$VmNetworkGateway = Get-SCVMNetworkGateway -Name $VmNetworkGatewayName -VMNetwork $VmNetwork

# Create a new Tenant Virtual Gateway if not configured
if($VmNetworkGateway -eq $null)
{
    if($EnableBGP -eq $false)
    {
        # Create a new Virtual Gateway for tenant
        $VmNetworkGateway = Add-SCVMNetworkGateway -Name $VmNetworkGatewayName -EnableBGP $false -
NetworkGateway $gatewayDevice -VMNetwork $VmNetwork -RoutingIPSubnet $GatewaySubnet;
    }
}

```

```

    else
    {
        if($TenantASN -eq "0")
        {
            Write-Verbose "Please specify valid ASN when using BGP"
            return
        }

        # Create a new Virtual Gateway for tenant
        $VmNetworkGateway = Add-SCVMNetworkGateway -Name $vmNetworkGatewayName -EnableBGP $true -NetworkGateway
$gatewayDevice -VMNetwork $vmNetwork -RoutingIPSubnet $GatewaySubnet -AutonomousSystemNumber $TenantASN;
    }

}

if ($VmNetworkGateway -eq $null)
{
    Write-Verbose "Could not Find / Create Virtual Gateway for $($VmNetwork.Name), quitting"
    return
}

# Check if the network connection already exists
$vpnConnection = Get-SCVPNConnection -VMNetworkGateway $VmNetworkGateway -Name $L3VPNConnectionName
if ($vpnConnection -ne $null)
{
    Write-Verbose "L3 Network Connection for $($VmNetwork.Name) already configured, skipping"
}
else
{
    # Create a new L3 Network connection for tenant
    $vpnConnection = Add-SCVPNConnection -NextHopNetwork $nexthopvmNetwork -Name $L3VPNConnectionName -
IPAddresses $LocalIPAddresses -PeerIPAddresses $PeerIPAddresses -VMNetworkGateway $VmNetworkGateway -protocol
L3;

    if ($vpnConnection -eq $null)
    {
        Write-Verbose "Could not add network connection for $($VmNetwork.Name), quitting"
        return
    }
    Write-Output "Created VPN Connection " $vpnConnection;
}

# Add all the required static routes to the newly created network connection interface
foreach($route in $RoutingSubnets)
{
    Add-SCNetworkRoute -IPSubnet $route -RunAsynchronously -VPNConnection $vpnConnection -VMNetworkGateway
$VmNetworkGateway
}
...

```

Configure L3 forwarding

L3 forwarding enables connectivity between the physical infrastructure in the datacenter and the virtualized infrastructure in the Hyper-V network virtualization cloud.

Using L3 forwarding connection, tenant network virtual machines can connect to a physical network through the Windows Server 2016/2019 SDN Gateway, which is already configured in an SDN environment. In this case, the SDN gateway acts as a router between the virtualized network and the physical network.

To learn more, check these articles: [Windows server gateway as a forwarding gateway](#) and [RAS gateway high availability](#).

Ensure the following before you attempt to configure L3 connection:

- Ensure you're logged on as an administrator on the VMM server.

- You must configure a unique next-hop logical network, with unique VLAN ID, for each Tenant VM network for which L3 forwarding needs to be set up. There must be 1:1 mapping between a tenant network and corresponding physical network (with unique VLAN ID).

Use the following steps to create the next-hop logical network in VMM:

1. On the VMM console, select **Logical Networks**, right-click, and select **Create Logical Network**.
2. In the **Settings** page, choose **One connected network** and select **Create a VM network with the same name to allow virtual machines to access this logical network directly** and **Managed by Microsoft Network Controller**.

Create Logical Network Wizard

Settings

Name

Settings

Network Site

Summary

Specify logical network settings

Select the option which describes this logical network:

☒ **One connected network**

The network sites within this network are equivalent and routable to one another and can be used as a single connected network.

☐ Allow new VM networks created on this logical network to use network virtualization

☒ Create a VM network with the same name to allow virtual machines to access this logical network directly

☒ Managed by Microsoft Network Controller

☐ Public IP address network

☐ **VLAN-based independent networks**

The subnet-VLAN pairs defined by the sites in this logical network are used as independent networks. They might or might not be routable to one another.

☐ **Private VLAN (PVLAN) networks**

The network sites within this logical network contain independent networks consisting of primary and secondary VLAN pairs in isolated mode.

Previous Next Cancel

3. Create an IP Pool for this new logical network. IP address from this pool is required for setting up L3 forwarding.

Use the following steps to configure L3 forwarding:

NOTE

You cannot limit bandwidth in L3 VPN connection.

1. In the VMM console, select the tenant virtual network that you want to connect to the physical network, through L3 gateway.
2. Right-click the selected tenant virtual network, select **Properties > Connectivity**.
3. Select **Connect to another network through a VPN tunnel**. Optionally, to enable BGP peering in your datacenter, select **Enable Border Gateway Protocol (BGP)**.

Tenant_Alpha_vNET Properties

Name

VM Subnets

Connectivity

VPN Connections

Border Gateway Protocol...

Access

Connectivity

By default, network virtualization provides an isolated network. Select the options below to add external connectivity to this network.

☒ Connect to another network through a VPN tunnel

☒ Enable Border Gateway Protocol (BGP)

☐ Connect directly to an additional logical network

☒ Direct routing

☐ Network address translation (NAT)

Select the gateway device to use for this connection.

Gateway Device: ContosoNCCluster

Device capabilities:

Site to Site VPN

Border Gateway Protocol (BGP)

Direct routing

Remaining capacity:

49 VM networks

250 VPN tunnels

View Script

OK Cancel

4. Select the network controller service for the gateway device.
5. In the **VPN Connections** page, click **Add > Add Layer 3 tunnel**.

Tenant_Alpha_vNET Properties

Name

VM Subnets

Connectivity

VPN Connections

Border Gateway Protocol...

Access

Routing Subnet

A subnet is required for routing packets out of the VM network.

Subnet:

Specify VPN connections:

+ Add - Remove

Add IPsec tunnel

Add GRE tunnel

Add Layer 3 tunnel

View Script

OK Cancel

6. Provide a subnet in the CIDR notation format for **Routing Subnet**. This subnet is used to route packets out of the VM network. You do not need to pre-configure this subnet in your datacenter.

Routing Subnet

A subnet is required for routing packets out of the VM network.

Subnet:

7. Use the following information and configure the L3 connection:

PARAMETER	DETAILS
Name	User-defined name for the L3 forwarding network connection.
VMNetwork (NextHop)	User-defined name for the next hop VM network, which was created as a prerequisite. This represents the physical network that wants to communicate with the tenant VM network. When you click <i>Browse</i> , only the <i>One Connected VM Networks</i> managed by Network service will be available for selection.
Peer IP Address	IP address of the physical network gateway, reachable over L3 logical network. This IP address must belong to the next hop logical network that you created as the prerequisite. This IP will serve as the next hop, once the traffic destined to the physical network from the tenant VM network reaches the SDN gateway. This must be an IPv4 address. There can be multiple peer IP addresses, must be separated by comma.
Local IP Addresses	IP addresses to be configured on the SDN gateway L3 network interface. These IP addresses must belong to the next hop logical network that you created as prerequisite. You must also provide the subnet mask. Example: 10.127.134.55/25. This must be an IPv4 address and should be in CIDR notation format. Peer IP address and Local IP addresses should be from the same Pool. These IP addresses should belong to the subnet defined in Logical Network Definition of VM Network.

- If you are using static routes, type all the remote subnets that you want to connect to, in **Routes**.

Tenant_Alpha_vNET Properties

Name

VM Subnets

Connectivity

VPN Connections

Border Gateway Protocol...

Access

Routing Subnet

A subnet is required for routing packets out of the VM network.

Subnet: 10.10.10.0/24

Specify VPN connections:

+ Add - Remove

Routes
0 routes

Subnet

Add

Remove

View Script

OK

Cancel

NOTE

You must configure routes in your physical network, for the tenant virtual network subnets, with the next hop as the IP address of the L3 interface on the SDN gateway (Local IP address used in the creation of L3 connection). This is to ensure that the return traffic to the tenant virtual network is routed correctly through the SDN gateway.

- If you are using BGP, ensure that BGP peering is established between the SDN gateway internal interface IP address, (which is present in a different compartment on the gateway VM, not the default compartment) and the peer device on the physical network.

For BGP to work, you must do the following steps:

1. Add BGP peer for the L3 connection. Enter your ASN, peer BGP IP, and its ASN on the **Border Gateway Protocol** page.

Tenant_Alpha_vNET Properties

Name

VM Subnets

Connectivity

VPN Connections

Border Gateway Protocol...

Access

Border Gateway Protocol (BGP)

BGP is a protocol that enables automatic exchanges of routes between individual networks. Enable BGP only if your environment is configured for it.

Autonomous system number (ASN):

BGP peers:

IP Address	Autonomous System Number
<input type="text"/>	<input type="text"/>

2. Determine the SDN gateway internal address as detailed in [the following section](#).
3. Create BGP peer on the remote end (physical network gateway). While creating the BGP peer, use the SDN gateway internal address (as determined in the previous step) as the peer IP address.
4. Configure a route on the physical network with the destination as the SDN gateway internal address and the next hop as the L3 interface IP address (Local IP address value used when creating L3 connection).

NOTE

After configuring the L3 connection, you must configure routes in your physical network for the tenant virtual network subnets, with the next hop as the IP address of the L3 interface on the SDN gateway (parameter `LocalIpAddresses` in the script). This is to ensure that the return traffic to the tenant virtual network is routed correctly through the SDN gateway.

You can choose to configure static routes or dynamic routes (over BGP) with the L3 connection. If you are using static routes, you can add them using **Add-SCNetworkRoute** as described in the script below.

If you use BGP with L3 tunnel connection, BGP peering must be established between the SDN gateway internal interface IP address, which is present in a different compartment on the gateway VM (not the default compartment) and the peer device on the physical network.

For BGP to work, you must do the following steps:

1. Add BGP peer for the L3 connection using the **Add-SCBGPPeer** cmdlet.

Example: `Add-SCBGPPeer -Name "peer1" -PeerIPAddress "12.13.14.15" -PeerASN 15 -VMNetworkGateway $VmNetworkGateway`

2. Determine the SDN gateway internal address as detailed in the [following section](#).
3. Create BGP peer on the remote end (physical network gateway). While creating the BGP peer, use the SDN gateway internal address (determined in Step 2 above) as the peer IP address.

4. Configure a route on the physical network with the destination as the SDN gateway internal address and the next hop as the L3 interface IP address (*LocalIPAddresses* parameter in the script).

Determine the SDN gateway internal address

Use the following procedure:

Run the following PowerShell cmdlets on a network controller installed computer or a computer that was configured as a network controller client:

```
$gateway = Get-NetworkControllerVirtualGateway -ConnectionUri <REST uri of your deployment>
$gateway.Properties.NetworkConnections.Properties.IPAddresses
```

The results of this command can display multiple virtual gateways, depending on how many tenants have configured gateway connections. Each virtual gateway can have multiple connections (IPSec, GRE, L3).

As you already know the L3 interface IP address (*LocalIPAddresses*) of the connection, you can identify the correct connection based on that IP address. After you have the correct network connection, run the following command (on the corresponding virtual gateway) to get the BGP router IP address of the virtual gateway

```
$gateway.Properties.BgpRouters.Properties.RouterIp
```

The result of this command provides the IP address that you must configure on the remote router as the peer IP Address.

Set up the traffic selector from VMM PowerShell

Use the following procedure:

NOTE

Values used are examples only.

1. Create the traffic selector by using the following parameters.

```
$t= new-object Microsoft.VirtualManager.Remoting.TrafficSelector

$t.Type=7 // IPV4=7, IPV6=8

$t.ProtocolId=6 // TCP =6, reference: https://en.wikipedia.org/wiki/List_of_IP_protocol_numbers

$t.PortEnd=5090

$t.PortStart=5080

$t.IpAddressStart=10.100.101.10

$t.IpAddressEnd=10.100.101.100
```

2. Configure the above traffic selector by using **-LocalTrafficSelectors** parameter of **Add-SCVPNConnection** or **Set-SCVPNConnection**.

Configure guest clusters in SDN through VMM

5/30/2019 • 2 minutes to read

This article explains about how to configure guest clusters in SDN using the System Center - Virtual machine Manager (VMM).

VMM was supporting guest clustering. However, with the advent of the network controller (NC), Windows Server 2016 and System Center 2016, the configuration of guest clusters has undergone some change.

With the introduction of the network controller, VMs, which are connected to the virtual network are only permitted to use the IP address that NC assigns for communication. NC does not support floating IP addresses which are essential for technologies such as Microsoft Failover Clustering to work.

VMM 1801 and later, you can enable this feature by emulating the floating IP functionality through the Software Load Balancer (SLB) in the SDN.

VMM 1801 and later supports guest clustering in SDN through an Internal Load Balancer(ILB) Virtual IP(VIP). The ILB uses probe ports which are created on the guest cluster VMs to identify the active node. At any given time, the probe port of only the active node responds to the ILB and all the traffic directed to the VIP is routed to the active node.

Before you start

Ensure the following prerequisite:

Guesting clustering is managed through the SDN NC. Ensure you have set up SDN and [deployed NC](#) and [SLB](#).

Procedure - configure guest clusters

Use the following steps:

1. Create a cluster for your VMs using the information provided in [this article](#).

NOTE

Cluster should have a ProbePort parameter set to a port of your choice. This port is required while configuring the VIP template in the next step. Note the reserved IP address you are using for this cluster. This IP will be required later while creating a VIP using the VIP template.

2. Create a VIP template. In VMM console > **Fabric** > **Networking** > **VIP Templates**, right-click and select **Add VIP Template**.

- In **Type**, select **Microsoft Network Controller**.

Load balancer VIP template Wizard

Type

Name

Type

Protocol

Load Balancing

Health Monitors

Summary

Specify a template type

Select the type of template to create that best matches your environment.

☐ Generic
A generic template can be used on any supported load balancer.

☒ Specific
A virtual IP template can be used only on a load balancer of a specific manufacturer and model.

Manufacturer: Microsoft

Model: Microsoft Network Controller

Previous Next Cancel

- In **Load Balancing**, select **Enable Floating IP** Checkbox.

Load balancer VIP template Wizard

Load Balancing

Name

Type

Protocol

Load Balancing

Health Monitors

Summary

Specify a load balancing method

Select one of the available load balancing methods.

Load balancing method: Default

Default uses 5 tuple (source IP, source port, protocol, destination IP and destination port) to create hash and distributes the traffic.

For guest clustering scenario, Enable Floating IP should be set.

☒ Enable Floating IP

Previous Next Cancel

- In **Health Monitor**, add the probe which would be used on the guest cluster VMs. Here, you must add the same port which was configured while clustering the hosts in the previous step.

Load balancer VIP template Wizard

Health Monitors

Name

Type

Protocol

Load Balancing

Health Monitors

Summary

Specify a health monitor

Specify one or more requests that you want to send at regular intervals to verify that a load balancer is available. Adding a health monitor is optional.

Health monitors:

Protocol	Request	Respo...	Port	Interval	Time...	Retries
▼				5	2	3

Insert

Delete

Previous Next Cancel

3. Using PowerShell, create a VIP using the VIP template.

NOTE

As explained in the beginning of this article, an Internal Load Balancer VIP is being implemented to support Guest Clustering. The PowerShell script for creating a VIP from the VIP template, to Load Balance Internal Network traffic is provided below.

Use the [sample script here](#) to create a VIP and load balance the internal network. Modify the parameters as required, based on the following details:

- LBServiceName – name of the Network Controller service
- VipMemberNames – names of the VMs in the cluster.
- VipNetworkName – name of the tenant network.
- VipAddress – the reserved IP address from the tenant network which was used in step 2 as the IP address for the VM cluster.
- VipTemplateName – name of the VIP template created in step 3, above.
- VipName – any friendly name you want to refer to the VIP by.

Update the network controller server certificate

5/30/2019 • 2 minutes to read

Network controller (NC) uses a certificate for Northbound communication with REST clients (such as VMM) and Southbound communication with Hyper-V hosts and software load balancers.

You can change or update this certificate in the following scenarios, after you deploy the NC.

- The certificate has expired
- You want to move from a self-signed certificate to a certificate that is issued by a certificate authority (CA).

NOTE

If you renew the existing certificate with the same key, these steps are not required.

Before you start

Make sure you create a new SSL certificate with existing network controller's REST name. [Learn more.](#)

Update the server certificate

1. If the certificate is self-signed, do the following:
 - Certificate with private key - Export the certificate and import it on all the NC nodes' **My** store.
 - Certificate without a private key - Export the certificate and import it on all the NC nodes' **Root** store.
2. If the certificate is a CA issued certificate, import it in all network controller nodes' **My** store.

NOTE

DO NOT remove the current certificate from the NC nodes. You should validate the updated certificate before you remove the existing one. Proceed with rest of the steps to update the certificate.

3. Update the server certificate by executing the following PowerShell command on one of the NC nodes.

```
$certificate = Get-ChildItem -Path Cert:\LocalMachine\My | Where {$_.Thumbprint -eq "Thumbprint of new certificate"}
Set-NetworkController -ServerCertificate $certificate
```

4. Update the certificate used for encrypting the credentials stored in the NC by executing the following command on one of the NC nodes.

```
$certificate = Get-ChildItem -Path Cert:\LocalMachine\My | Where {$_.Thumbprint -eq "Thumbprint of new certificate"}
Set-NetworkControllerCluster -CredentialEncryptionCertificate $certificate
```

5. Retrieve a server REST resource by executing the following PowerShell command on one of the NC nodes.

```
Get-NetworkControllerServer -ConnectionUri <REST uri of your deployment>
```

6. In the Server REST resource, navigate to the **Credentials** object and check the credential of type **X509Certificate** with a value matching your certificate's thumbprint. Note the credential resource ID.

```
"Connections":
{
  {
    "ManagementAddresses": [ "contoso.com" ],
    "CredentialType": "X509Certificate",
    "Protocol": null,
    "Port": null,
    "Credential": {
      "Tags": null,
      "ResourceRef": "/credentials/<credential resource Id>",
      "InstanceId": "00000000-0000-0000-0000-000000000000",
      ...
    }
  }
}
```

7. Update the credential REST resource of type **X509Certificate** retrieved above with the thumbprint of the new certificate.

Execute these PowerShell cmdlet on any of the NC Node.

```
$cred=New-Object Microsoft.Windows.Networkcontroller.credentialproperties
$cred.type="X509Certificate"
$cred.username=""
$cred.value="<thumbprint of the new certificate>"
New-NetworkControllerCredential -ConnectionUri <REST uri of the deployment> -ResourceId <credential
resource Id> -Properties
$cred
```

8. If the new certificate is a self-signed certificate, provision the certificate (without the private key) in the trusted root certificate store of all the Hyper-V hosts and software load balancer MUX virtual machines.
9. Provision the NC certificate (without the private key) in the trusted root certificate store of the VMM machine using the following PowerShell cmdlet:

```
$certificate = Get-SCCertificate -ComputerName "NCRestName"
$networkservice = Get-SCNetworkService | Where {$_.IsNetworkController -eq $true}
Set-SCNetworkService -ProvisionSelfSignedCertificatesforNetworkService $true -Certificate
$certificate -NetworkService $networkservice
```

- **NetworkService** is the network controller service, **Certificate** is the new NC server certificate.
 - **ProvisionSelfSignedCertificatesforNetworkService** is **\$true** if you are updating to a self-signed certificate.
10. Verify that the connectivity is working fine with the updated certificate.

You can now remove the previous certificate from the NC nodes.

Next steps

Validate the [NC deployment](#) to ensure that the deployment is successful.

Configure SLB VIPs through VMM service templates

5/30/2019 • 3 minutes to read

Software Defined Networking (SDN) in Windows 2016 can use Software Load Balancing (SLB) to evenly distribute network traffic among workloads managed by service provider and tenants. VMM 2016 only supports deploying SLB Virtual IPs (VIPs) using power shell.

System Center - Virtual Machine Manager (VMM) 1801 and later supports configuration of SLB VIPs while deploying multi-tier application by using service templates and also supports both public and internal load balancing.

Before you begin

Ensure the following prerequisites are met:

- [Deployed SDN Network Controller](#).
- [Deployed SDN Software load balancer](#).

Procedure to create SLB VIPs

Use the following steps:

1. Specify the affinity to logical networks.

- In the VMM console, click **Fabric > Network Service > Network Controller > Properties > Logical Network Affinity** page.
- Specify the Front-end and Back-end networks available for load balancing and click **OK**.

NC Properties

Logical Network Affinity

Specify the affinity to logical networks

The virtual IP (VIP) address that is assigned to a load balancer is the front-end address for a service. The dedicated IP addresses that are assigned to the virtual machines in a service tier are the back-end addresses. Specify from which logical networks VMM should assign the front-end and back-end IP addresses. If you do not select at least one front-end and back-end logical network affinity, the load balancer will not be selected during placement.

Logical networks:

Logical Network	Front-end	Back-end
HNVPA	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
PrivateVIP	<input type="checkbox"/>	<input type="checkbox"/>
PublicVIP	<input type="checkbox"/>	<input type="checkbox"/>
Transit	<input type="checkbox"/>	<input type="checkbox"/>

View Script OK Cancel

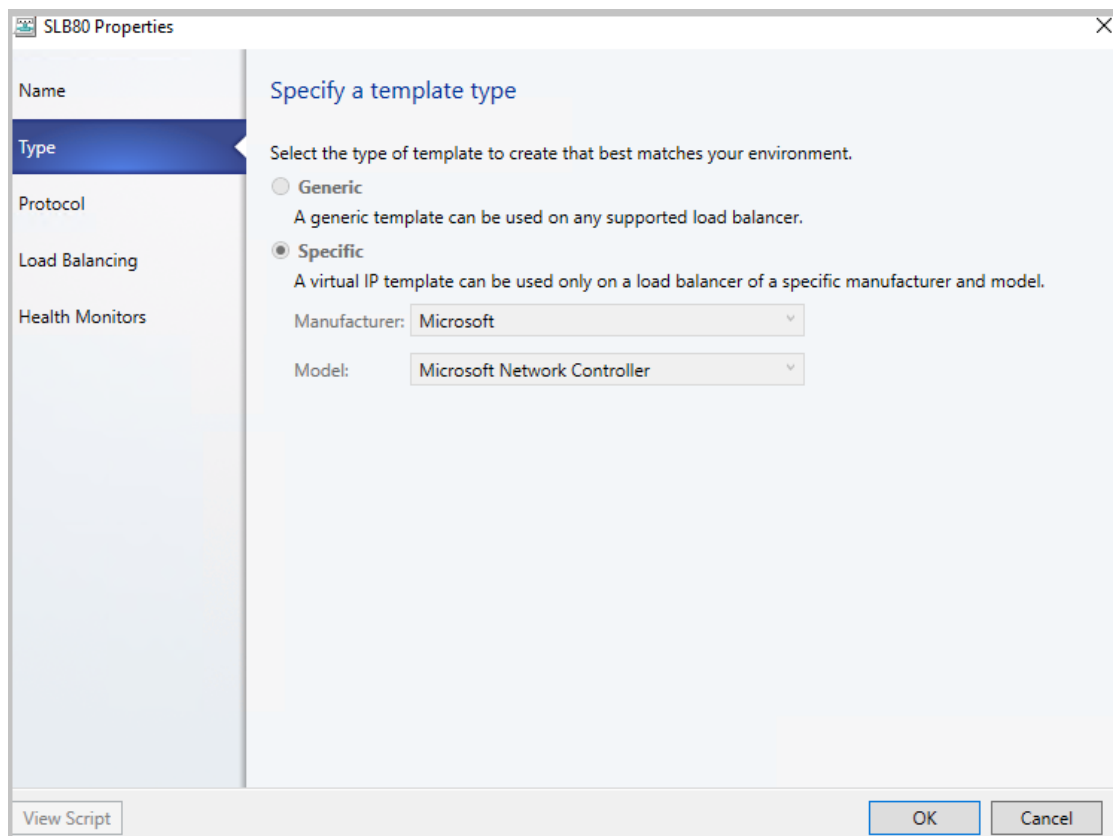
2. Create a VIP template.

- In the VMM console click **Fabric > Create VIP Template**.
- In the **Load Balancer VIP Template Wizard > Name**, specify the template name and description.
- In **Virtual IP port**, specify the port that will be used for the type of network traffic you want to balance.
- In **Backend port**, specify the port on which the backend server is listening for requests.

The screenshot shows the 'SLB80 Properties' dialog box with the 'Name' tab selected. The dialog has a left sidebar with tabs: Name, Type, Protocol, Load Balancing, and Health Monitors. The main area is titled 'Specify a name and description for this template' and contains a text box for 'Template name' (containing 'SLB80'), a larger text box for 'Description', and two numeric input boxes for 'Virtual IP port' and 'Backend port' (both containing '80'). At the bottom are buttons for 'View Script', 'OK', and 'Cancel'.

Property	Value
Template name	SLB80
Description	
Virtual IP port	80
Backend port	80

- In **Type**, specify a template type, click **Specific**. Select **Microsoft** for Manufacturer. Select **Microsoft network controller** for Model. Click **Next**.



The image shows the 'SLB80 Properties' dialog box with the 'Type' tab selected. The left sidebar contains 'Name', 'Type', 'Protocol', 'Load Balancing', and 'Health Monitors'. The main area is titled 'Specify a template type' and contains instructions to select a template type. Two radio buttons are present: 'Generic' and 'Specific'. The 'Specific' option is selected. Below the radio buttons, there are two dropdown menus: 'Manufacturer' set to 'Microsoft' and 'Model' set to 'Microsoft Network Controller'. At the bottom, there is a 'View Script' button on the left and 'OK' and 'Cancel' buttons on the right.

SLB80 Properties

Name

Type

Protocol

Load Balancing

Health Monitors

Specify a template type

Select the type of template to create that best matches your environment.

☐ Generic
A generic template can be used on any supported load balancer.

☒ Specific
A virtual IP template can be used only on a load balancer of a specific manufacturer and model.

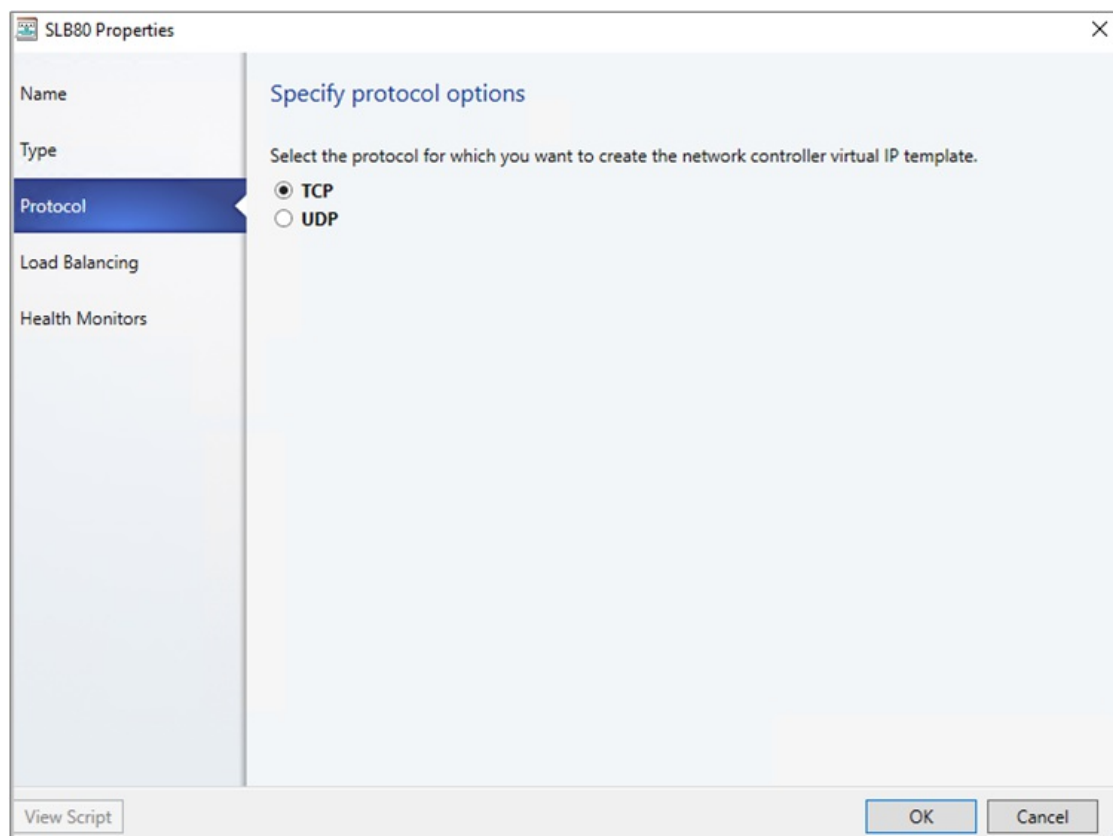
Manufacturer: Microsoft

Model: Microsoft Network Controller

View Script

OK Cancel

- In **Protocol**, specify protocol options. Click **Next**.



The image shows the 'SLB80 Properties' dialog box with the 'Protocol' tab selected. The left sidebar contains 'Name', 'Type', 'Protocol', 'Load Balancing', and 'Health Monitors'. The main area is titled 'Specify protocol options' and contains instructions to select a protocol. Two radio buttons are present: 'TCP' and 'UDP'. The 'TCP' option is selected. At the bottom, there is a 'View Script' button on the left and 'OK' and 'Cancel' buttons on the right.

SLB80 Properties

Name

Type

Protocol

Load Balancing

Health Monitors

Specify protocol options

Select the protocol for which you want to create the network controller virtual IP template.

☒ TCP

☐ UDP

View Script

OK Cancel

- In **Load Balancing method**, select the method and click **Next**.

SLB80 Properties

Name

Type

Protocol

Load Balancing

Health Monitors

Specify a load balancing method

Select one of the available load balancing methods.

Load balancing method: **Default**

Default uses 5 tuple (source destination port) to create h... IP and

For guest clustering scenario, Enable Floating IP should be set.

☐ Enable Floating IP

View Script OK Cancel

- In **Health Monitors**, you can optionally specify that a verification should run against the load balancer at regular intervals. To add a health monitor, specify the protocol and the request. For example, entering the command GET makes an HTTP GET request for the home page of the load balancer and checks for a header response. You can also modify the response type, and monitoring interval, timeout, and retries. Note that the timeout should be less than the interval.

SLB80 Properties

Name

Type

Protocol

Load Balancing

Health Monitors

Specify a health monitor

Specify one or more requests that you want to send at regular intervals to verify that a load balancer is available. Adding a health monitor is optional.

Health monitors:

Protocol	Request	Respo...	Port	Interval	Time...	Retries
[v]				5	2	3

Http
Tcp

Insert
Delete

View Script OK Cancel

- In **Summary**, confirm the Settings and click **Finish** to create the VIP template.

3. Configure SLB VIP while deploying Service

- If the service template isn't open, click **Library > Templates > Service Templates** and open it.
- Click **Actions > Open Designer**.
- In the **Service Template Designer**, click the **Service Template Components group > Add Load Balancer**.
- Click the load balancer object. You'll identify it with the VIP template name.
- Click **Tool > Connector**. Click the Server connection associated with template and then click a NIC object to connect the load balancer to the adapter. In the **NIC properties**, check the address types and that the MAC address is static.

Note that the server connection must be connected to the Back-End network interface of the service. The back-end network interface can be connected to either a One Connected VM Network or a network virtualized VM Network.

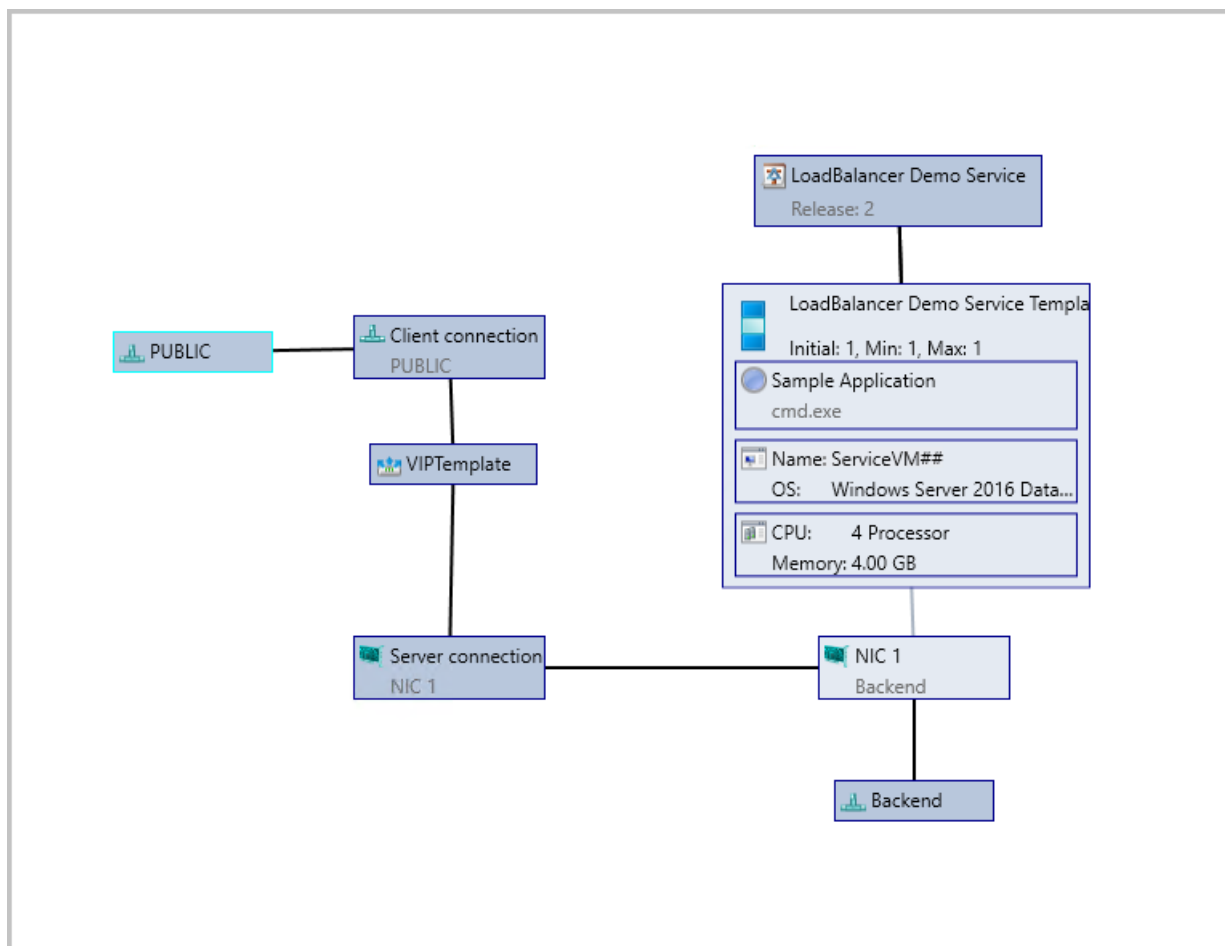
- With the Connector enabled, click the client connection associated with the load balancer and then click a logical network object.

NOTE

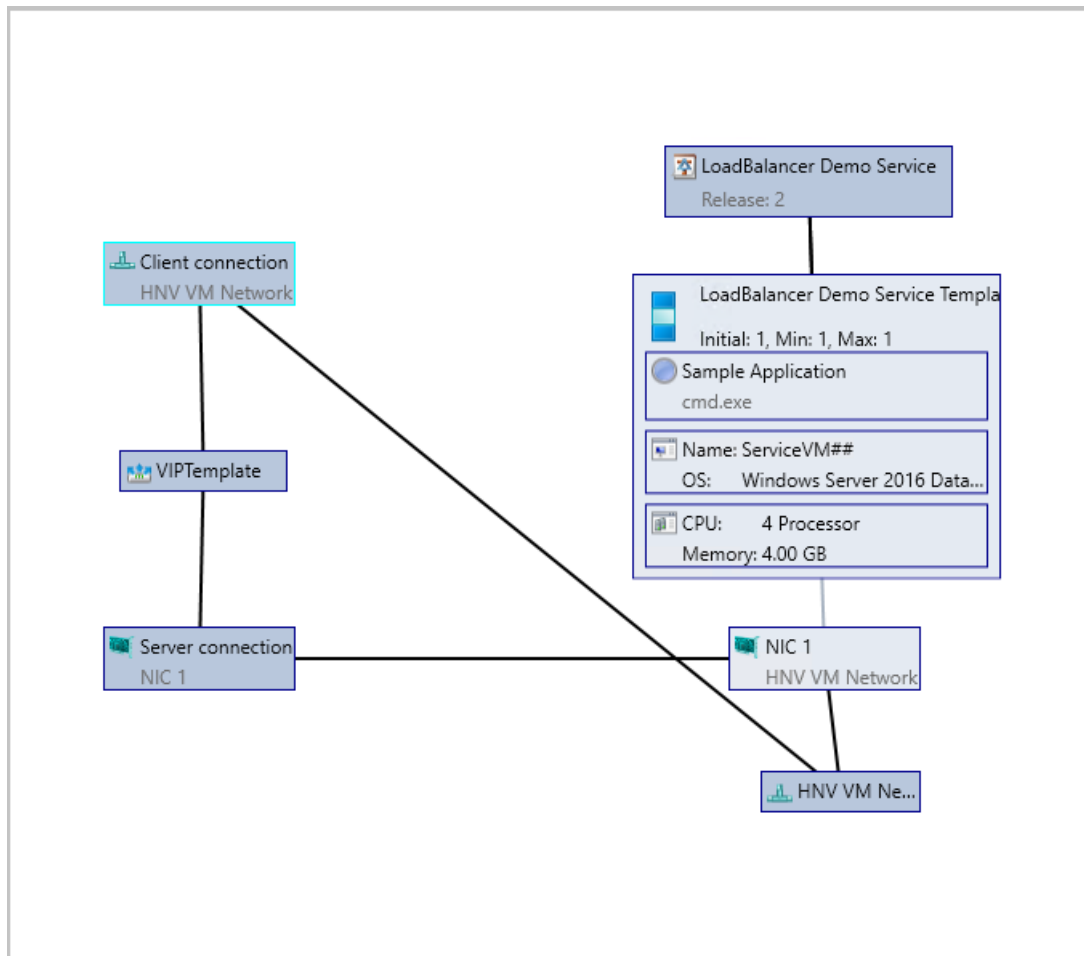
Client connection must be connected to a Front-End network of the load balancer. This can be a Public VM network or a network virtualized VM network. A network virtualized VM Network is used for internal load balancing scenarios.

- Save the service template in **Service Template > Save and Validate**.

Example 1: Configuring Service with 'Public' VM Network as front end. Here the 'Backend' network can be One connected or network virtualized VM network.



Example 2: Configuring Service with Front-end and Back-end connected to network virtualized VM Network 'HNV VM Network'. This scenario is used for Internal load balancing.



Set up the VIP for user access

When the service is deployed, VMM automatically selects a VIP from the reserved range in the static IP address pool, and assigns it to the load-balanced service tier. To enable users to connect to the service, after the service is deployed you need to determine the VIP and configure a DNS entry for it.

1. After the service is deployed click **Fabric > Networking > Load Balancers**.
2. Click **Show > Service > Load Balancer Information for Services** and expand the service to see which VIP is assigned.
3. If users use the DNS name to access the service, request that the DNS administrator manually create a DNS entry for the VIP. The entry should be the name that users will specify to connect to the service. For example, servicename.contosol.com.

Back up and restore the SDN infrastructure

8/29/2019 • 4 minutes to read

This article describes the backup and recovery process of a software defined network (SDN) infrastructure in the System Center Virtual Machine Manager (VMM) fabric, and provides any applicable recommendations for the process completion.

To back up and restore an SDN, you must back up and restore the network controller (NC) that is deployed in the SDN. Use the following procedures in the sequence explained.

Back up the network controller

Back up the network controller database by using the network controller Rest API. [Learn more.](#)

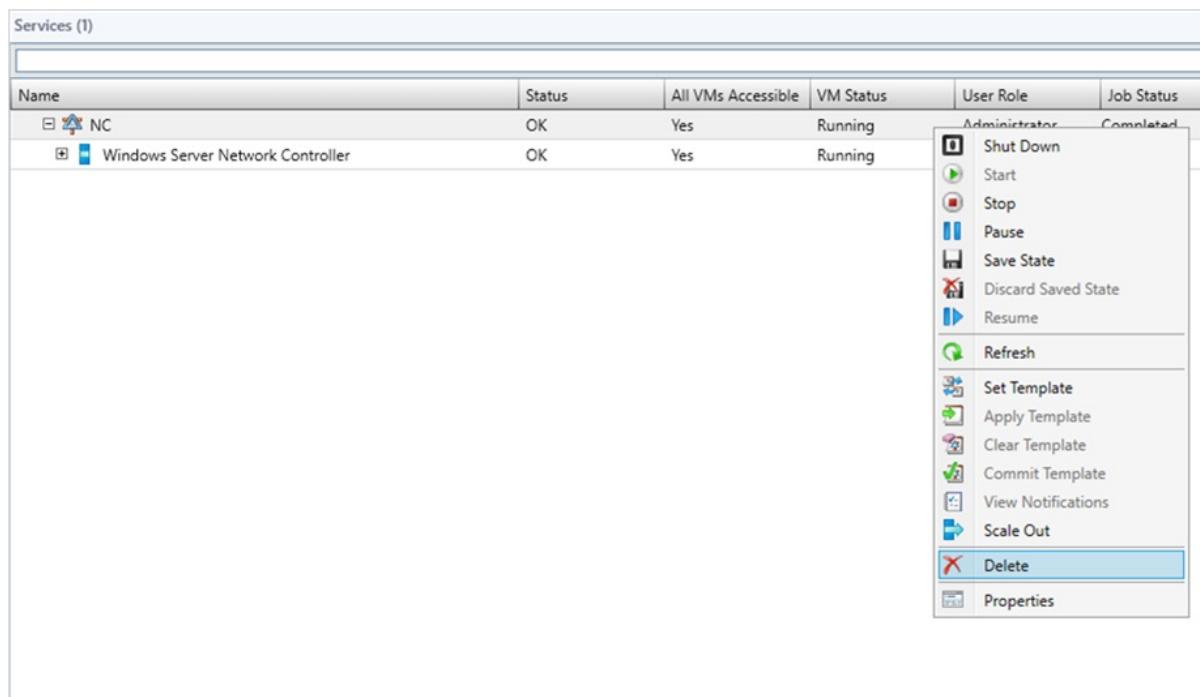
Bring up the new network controller

Use the following procedures to bring up a new network controller:

1. In the VMM console, **VMs and Services** > **Services**, select the network controller service instance and click **Delete**.

NOTE

Remove the network controller service instance only. Do not remove the network controller from **Fabric > Network Services**.



2. Ensure that the DNS record for the network controller is removed from the DNS.
3. Deploy a new network controller service instance from the VMM by using the same service deployment settings that were used for the original service instance deployment. [Learn more.](#)
4. Verify that the deployment job is successful.

Restore the network controller from a backup

Restore the network controller from a network controller backup by using the network controller Rest API. [Learn more](#).

Refresh the network controller and synchronize VMM and NC

Depending on the SDN state captured in the network controller backup and the current VMM state, some of the resources in VMM and network controller might be out of sync.

Use the following refresh procedures to find any such differences between VMM and NC, and accordingly resolve them.

NOTE

- Refresh cmdlets for refreshing network controller objects are available from VMM 2016 UR3.
- If the network controller contains any objects which are not present in the VMM DB, then the VMM will not refresh (even if those objects are created by using VMM earlier). Delete those objects from NC and recreate the objects from VMM to manage these objects from VMM again.

Refresh port ACLs

1. Get all the NC-managed port ACLs from the VMM server by using the following cmdlet:

```
$portACLs = Get-SCPortACL | Where-Object {$_.ManagedByNC -eq $True}
```

2. Run the **Read-SCPortACL** cmdlet on all the NC- managed port ACLs to refresh.

```
foreach($portACL in $portACLs)
{
    Read-SCPortACL -PortACL $portACL
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh logical networks

1. Get all the NC-managed logical networks from the VMM server by using the following cmdlet:

```
$logicalNetworks = Get-SCLogicalNetwork | Where-Object {$_.IsManagedByNetworkController -eq $True}
```

2. Run the **Read-SCLogicalNetwork** cmdlet on all the NC-managed logical networks to refresh.

```
foreach($logicalNetwork in $logicalNetworks)
{
    Read-SCLogicalNetwork -LogicalNetwork $logicalNetwork
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh gateways and load balancer muxes

1. Get all the gateways and load balancer muxes by using the following cmdlet:

```
$networkService = Get-SCNetworkService | Where-Object {$_.Model -eq 'Microsoft Network Controller'}
$fabricRoles = Get-SCFabricRole -NetworkService $networkService
$fabricRoleResources = @()
foreach($fabricRole in $fabricRoles)
{
    $fabricRoleResources += $fabricRole.ServiceVMs
}
$fabricRoleResources
```

2. Run the **Read-SCFabricRoleResource** cmdlet to refresh.

```
foreach($fabricRoleResource in $fabricRoleResources)
{
    Read-SCFabricRoleResource -FabricResource $fabricRoleResource
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh NAT connections and NAT rules

1. Get all the NAT connections by using the following cmdlet:

```
$vmNetworks = Get-SCVMNetwork | Where-Object {$_.NetworkManager.Model -eq 'Microsoft Network Controller'
-and $_.IsolationType -eq 'WindowsNetworkVirtualization'}

$natConnections = @()
foreach($vmNetwork in $vmNetworks)
{
    $natConnections += $vmNetwork.NATConnections
}
$natConnections
```

2. Run the **Read-SCNATConnection** cmdlet to refresh NAT connections and NAT rules.

```
foreach($natConnection in $natConnections)
{
    Read-SCNATConnection -NATConnection $natConnection
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh all load balancer VIPs

1. Get all the load balancer VIPs configured on NC by using the following cmdlet:

```
$loadBalancerVIPs = Get-SCLoadBalancerVIP | Where-Object {$_.LoadBalancer.Model -eq 'Microsoft Network Controller'}
```

2. Run the **Read-SCLoadBalancerVIP** cmdlet to refresh all the load balancer VIPs.

```
foreach($loadBalancerVIP in $loadBalancerVIPs)
{
    Read-SCLoadBalancerVIP -LoadBalancerVIP $loadBalancerVIP
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh VM Networks

1. Get all the NC-managed HNV VM networks from the VMM server by using the following cmdlet:

```
$VMNetworks = Get-SCVMNetwork | Where-Object {$_.NetworkManager.Model -eq 'Microsoft Network Controller'
-and $_.IsolationType -eq 'WindowsNetworkVirtualization'}
```

2. Run the **Read-SCVMNetwork** cmdlet on all the VM networks to refresh.

```
foreach($VMNetwork in $VMNetworks)
{
    Read-SCVMNetwork -VMNetwork $VMNetwork
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh gateway pools

1. Get the gateway fabric role from the VMM server by using the following cmdlet:

```
$networkService = Get-SCNetworkService | Where-Object {$_.Model -eq 'Microsoft Network Controller'}
$gatewayFabricRole = Get-SCFabricRole -NetworkService $networkService | Where-Object {$_. RoleType -eq
'Gateway '}
```

2. Run the **Read-SCFabricRole** cmdlet to refresh the fabric role.

```
foreach($fabricRole in $gatewayFabricRole )
{
    Read-SCFabricRole -FabricRole $fabricRole
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Refresh VM network gateways

1. Get all the VM network gateways that are configured for the VM networks by using the following cmdlet:

```
$vmNetworks = Get-SCVMNetwork | Where-Object {$_.NetworkManager.Model -eq 'Microsoft Network  
Controller' -and $_.IsolationType -eq 'WindowsNetworkVirtualization' -and $_.VMNetworkGateways.Count -  
gt 0}}  
$VMNetworkGateways = @()  
foreach($vmNetwork in $vmNetworks)  
{  
    $VMNetworkGateways += $vmNetwork.$VMNetworkGateways  
}
```

2. Run the **Read-SCVMNetworkGateway** cmdlet to refresh the gateways.

```
foreach($VMNetworkGateway in $VMNetworkGateways)  
{  
    Read-SCVMNetworkGateway -VMNetworkGateway $VMNetworkGateway  
}
```

3. Verify the VMM jobs' log for the result status and follow the recommendations from the log, in case of any failures.

Remove a Software Defined Network (SDN) from VMM fabric

12/7/2018 • 2 minutes to read

To remove a SDN from the System Center Virtual Machine Manager (VMM) fabric, you must remove the following objects - in the specified order:

- VM networks (associated with the NC managed logical networks).
- Logical networks (managed by NC).
- Software load balancer (if deployed or deployed and configured).
- Gateway (if deployed or deployed and configured).
- Network controller (if deployed or deployed and configured).

Remove the VM networks

NOTE

Ensure that no VMs or NICs are connected to the VM networks that you want to remove.

1. Click **VMs and Services > VM Networks**, select the VM network to remove.
2. Right-click the VM network and click **Delete**.
3. Repeat steps 1&2 for each VM network that you need to remove.

Remove the logical networks

NOTE

Ensure that no port profiles are associated with the logical networks that you want to remove.

1. Click **Fabric > Logical networks**, select the logical network to remove.
2. Right-click the logical network and click **Remove**.
3. Repeat steps 1&2 for each logical network that you need to remove.

NOTE

Logical networks associated with the SLB cannot be removed from the console. Use force delete to remove these (by using - **Force** flag).

Remove the software load balancer

1. Click **Fabric > Network Services**, select the software load balancer role.
2. Under **Services > Associated Services**, click **Browse** and then click **Clear Selection**.

This action removes the software load balancer service. Ensure the job is complete. If the job fails, restart the job after making the required changes that the error message details you.

3. Uncheck the pools that are associated with the SLB, except for the private VIP pool that is associated with the SLB Manager VIP.
4. To complete the removal of the SLB, force delete the private VIP pool, corresponding logical network definition and logical networks ((by using **-Force** flag).

Remove the gateway

1. Click **Fabric > Network Services**, select the Gateway manager role.
2. Under **Services > Associated Services**, click **Browse** and then click **Clear Selection**.

This action removes the gateway service. Ensure the job is complete. If the job fails, restart the job after making the required changes that the error message details you.

3. To complete the removal of the gateway, remove the gateway pool by using the following PowerShell scripts:

```
$nc=get-scnetworkservice | Where {$_.Model -eq "Microsoft Network Controller"}
$gwrole=get-scfabricrole -NetworkService $nc | Where {$_.RoleType -eq "Gateway"}
Set-SCFabricRole -FabricRole $gwrole -GatewayConfiguration $null
```

Remove the network controller

NOTE

Ensure that SLB/GW and associated logical networks are successfully removed.

1. Click **Fabric > Network Services**, select the network controller.
2. Right-click the NC and click **Remove**.

This action removes the NC service. Ensure the job is complete. If the job fails, restart the job after making the required changes that the error message details you.

Manage SDN resources in the VMM fabric

12/7/2018 • 3 minutes to read

This article summarizes the software-defined network (SDN) operations that you can manage in the System Center Virtual Machine Manager (VMM) fabric. For operations that can't be managed in the fabric, you need to use REST APIs or Windows Server PowerShell.

A software-defined network (SDN) abstract physical hardware network infrastructure into virtual networks. In the VMM fabric you can deploy and manage an SDN infrastructure, including network controller, software load balancers, and gateways, to provision and manage virtual networks at scale. [Learn more](#).

What can I manage in VMM?

SDN resources fall into two broad categories in VMM:

- **Known resources:** Resources that can be created and managed with VMM.
- **Unknown resources:** Resources that must be created and managed outside of VMM.

Known resources

These resources can be created and managed with or without VMM. If you make changes to these resources outside VMM, VMM overwrites the out-of-box changes, when a VMM operation is performed on the object. This could cause configuration and connectivity issues, and should be avoided whenever possible. There's no way to revert an overwrite unless you detect the issue, and reconfigure manually.

We strongly recommend that you configure resources that are known to VMM in the VMM fabric only.

KNOWN OBJECT	DETAILS	MODIFY
AccessControlList	An AccessControlList contains a list of ACL rules, and can be assigned to virtual subnets or IP configurations.	Overwritten by VMM if you enable out-of-box
AclRule	Summarizes the network traffic that is allowed or denied for a VM network interface.	Overwritten by VMM if you enable out-of-box
Gateway	Provides gateway services to one or more virtualNetworks.	Overwritten by VMM if you enable out-of-box
GatewayPool	GatewayPools aggregate a set of gateways resources into a single pool.	Overwritten by VMM if you enable out-of-box
Host		Overwritten by VMM if you enable out-of-box
HostProperties		Overwritten by VMM if you enable out-of-box
IpConfigurations	IP addresses of the load balancer	Overwritten by VMM if you enable out-of-box

KNOWN OBJECT	DETAILS	MODIFY
IpPool	Create an IP address pool on the network controller	Overwritten by VMM if you enable out-of-box
LoadBalancerManager	Configures the load balancing service of the Network Controller.	Overwritten by VMM if you enable out-of-box
LoadBalancerMux	Represents a MUX VM deployed in the network controller fabric.	Overwritten by VMM if you enable out-of-box
LogicalSubnets	A subnet/VLAN pair.	Overwritten by VMM if you enable out-of-box
MACPool	Creates a MAC address pool on the network controller	Overwritten by VMM if you enable out-of-box
NatRules	Configures the load balancer to apply NAT to traffic	Overwritten by VMM if you enable out-of-box
NetworkInterface	Specifies the configuration of either a host virtual interface (host vNIC) or a virtual server NIC (VMNIC).	Overwritten by VMM if you enable out-of-box
PortSettings		Overwritten by VMM if you enable out-of-box
PublicIpAddress	Specifies an IP address, which is publically available. It's used by virtualGateways and loadBalancers to indicate the IP address that can be used to communicate with the virtual network from outside.	Overwritten by VMM if you enable out-of-box
QualityOfService		Overwritten by VMM if you enable out-of-box
Servers	Represents a physical server that is being controlled by the Network Controller.	Overwritten by VMM if you enable out-of-box
VirtualGateway	Describes the gateway used for cross-premises connectivity from the virtual network.	Overwritten by VMM if you enable out-of-box
VirtualGatewayBgpPeer	Configures BGP peers of the virtualGateways resource.	Overwritten by VMM if you enable out-of-box
VirtualNetwork	Used to create a virtual network using HNV for tenant overlays.	Overwritten by VMM if you enable out-of-box
VirtualServer	Corresponds to a virtual machine. Must be created for VMs that correspond to gateway and MUX resources.	Overwritten by VMM if you enable out-of-box
VirtualSubnet	Used to create virtual subnets (VSIDs) under a tenant's virtual network (RDID).	Overwritten by VMM if you enable out-of-box

KNOWN OBJECT	DETAILS	MODIFY
VirtualSwitchManager	Configures the virtual switch properties on every server managed by the Network Controller	Overwritten by VMM if you enable out-of-box
VM	Corresponds to a virtual machine.	Overwritten by VMM if you enable out-of-box

Unknown resources

These resources to be created and managed outside the VMM fabric. VMM has no knowledge of them, and obviously doesn't overwrite them when they're configured outside the VMM console.

Unknown objects are any Network Controller resources that aren't listed in the table above. Get the [latest list](#) of SDN resources.

Manage Storage Spaces Direct in VMM

12/7/2018 • 2 minutes to read

This article provides an overview of Storage Spaces Direct (S2D), and how it's deployed in the System Center - Virtual Machine Manager (VMM) fabric.

Storage Spaces Direct (S2D) was introduced in Windows Server 2016. It groups physical storage drives into virtual storage pools, to provide virtualized storage. With virtualized storage you can:

- Manage multiple physical storage sources as a single virtual entity.
- Get inexpensive storage, with and without external storage devices.
- Gather different types of storage into a single virtual storage pool
- Easily provision storage, and expand virtualized storage on demand by adding new drives.

How does it work?

S2D creates pools of storage from storage that's attached to specific nodes in a Windows Server cluster. The storage can be internal on the node, or disk devices that are directly attached to a single node. Supported storage drives include NVMe, SSD connected via SATA or SAS, and HDD. [Learn more](#).

- When you enable S2D on a Windows Server cluster, S2D automatically discovers eligible storage, and adds it to a storage pool for the cluster.
- S2D also creates a built-in server-side storage cache to maximize performance. The fastest drives are used for caching, and the remaining drives for capacity. [Learn more](#) about the cache.
- You create volumes from a storage pool. Creating a volume creates the virtual disk (storage space), partitions and formats it, adds it to the cluster, and converts it to a clustered shared volumes (CSV).
- You configure different levels of fault tolerance for a volume, to specify how virtual disks are spread across physical disks in the pool, using SMB 3.0. You can configure a volume with no resiliency, or with mirror or parity resilience. [Learn more](#).

Converged and non-converged deployment

A cluster running S2D can be deployed in a couple of ways:

- **Hyper-converged deployment:** Hyper-V compute and S2D storage run within the same cluster, with no separation between them. This provides simultaneous scaling of compute and storage resources.
- **Disaggregated deployment:** Compute resources run on one Hyper-V cluster. S2D storage runs on a different cluster. You scale the clusters separately for finely tuned management.

Hyper-converged deployment

Here's an illustration for hyper-converged deployment

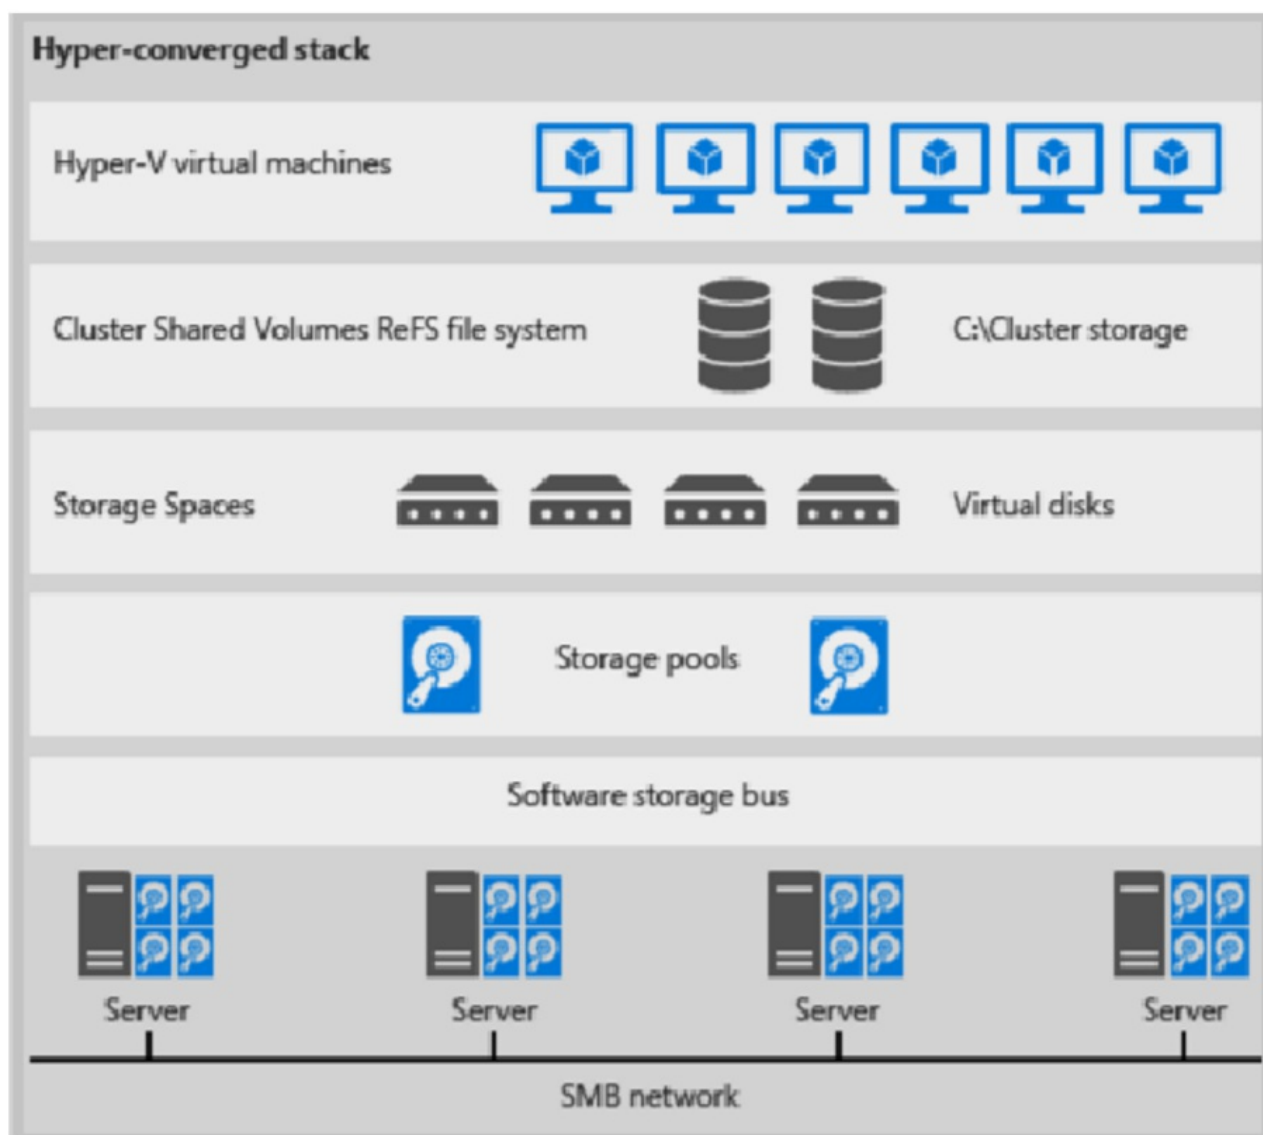


Figure 1: Hyper-converged deployment

- VM files are stored on local CSVs.
- File shares and SMB aren't used.
- After S2D CSV volumes are available, you provision them as you would any other Hyper-V deployment.
- You scale the Hyper-V compute cluster together with its S2D storage.

Disaggregated deployment

Here's an illustration for disaggregated deployment




Figure 2: Disaggregated deployment

- File shares are created on the S2D CSVs.
- Hyper-V VMs are configured to store their files on the scaled-out file server (SOFS), and accessed using SMB 3.0.
- You scale the Hyper-V and SOFS clusters separately for finely tuned management. For example, compute nodes might be near full capacity for the number of VMs, but storage nodes might have excess disk and IOPS capacity, so you add only additional compute nodes.

Next steps

[Deploy a hyper-converged S2D cluster](#)

Deploy a Storage Spaces Direct hyper-converged cluster in VMM

3/28/2019 • 4 minutes to read

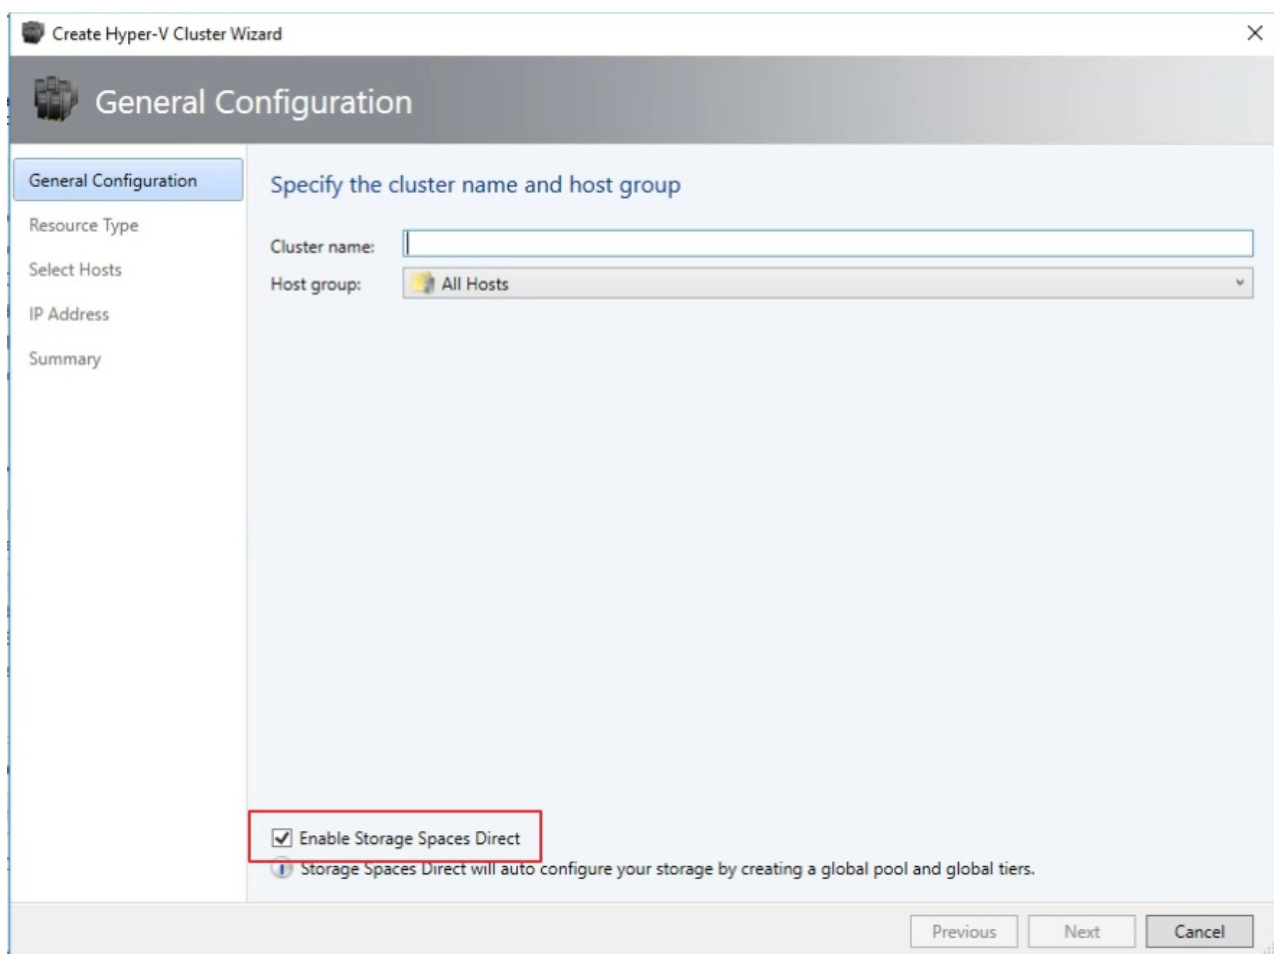
Read this article to set up a hyper-converged cluster running Storage Spaces Direct (S2D) in System Center - Virtual Machine Manager (VMM). [Learn more](#) about S2D.

You can deploy a hyper-converged S2D cluster by provisioning a Hyper-V cluster and enable S2D from existing Hyper-V hosts, or by provisioning from bare-metal servers.

You can't currently enable S2D in a hyper-converged deployment on a Hyper-V cluster deployed from bare metal computers with the Nano operating system.

NOTE

You must enable S2D before adding the storage provider to VMM. To enable S2D, go to **General Configuration > Specify the cluster name and host group** and select the **Enable Storage Spaces Direct** option, as shown below:



After you enable a cluster with S2D, VMM does the following:

1. The File Server role and the Failover Clustering feature is enabled.
2. Storage replica and data deduplication is enabled.
3. The cluster is optionally validated and created.
4. S2D is enabled, and a storage array is created with the same name you provide in the wizard.

If you use PowerShell to create a hyper-converged cluster, the pool and the storage tier is automatically created with the **Enable-ClusterS2D autoconfig=true** option.

Before you start

- Make sure you're running VMM 2016.
- Hyper-V hosts in a cluster should be running Windows Server 2016 with the Hyper-V Role installed, and be configured to host VMs.

After these prerequisites are in place, you provision a cluster, and set up storage resources on it. You can then deploy VMs on the cluster, or export the storage to other resources using SOFS.

Step 1: Provision the cluster

Provision a cluster from Hyper-V hosts

1. If you need to add the Hyper-V hosts to the VMM fabric, [follow these steps](#). If they're already in the VMM fabric, skip to the next step.
2. Follow the instructions for [provisioning a cluster from standalone Hyper-V hosts managed in the VMM fabric](#).

Note the following:

- When you set up the cluster, remember to select the **Enable Storage Spaces Direct** option on the **General Configuration** page of the Create Hyper-V Cluster wizard. -In **Resource Type**, select **Existing servers running a Windows Server operating system**, and select the Hyper-V hosts to add to the cluster.
- If S2D is enabled, you must validate the cluster. Skipping this step isn't supported.

Provision a cluster from bare metal machines

NOTE

Typically, S2D node requires RDMA, QOS and SET settings. To configure these settings for a node using bare metal computers, you can use the post deployment script capability in PCP. Here is the [sample PCP post deployment script](#). You can also use this script to configure RDMA, QoS and SET while adding a new node to an existing S2D deployment from bare metal computers.

1. Read the [prerequisites](#) for bare-metal cluster deployment. Note that:
 - The generalized VHD or VHDX in the VMM library should be running Windows Server 2016 with the latest updates. The **Operating system** and **Virtualization platform** values for the hard disk should be set.
 - For bare-metal deployment you need to add a pre-boot execution environment (PXE) server to the VMM fabric. The PXE server is provided through Windows Deployment Services. VMM uses it's own WinPE image, and you need to make sure that it's the latest. To do this, click **Fabric > Infrastructure > Update WinPE image**, and make sure that the job finishes.
2. Follow the instructions for [provisioning a cluster from bare-metal computers](#).

Step 2: Set up networking for the cluster

After the cluster is provisioned and managed in the VMM fabric, you need to set up networking for cluster nodes.

1. Start by [creating a logical network](#) to mirror your physical management network.
2. You need to [set up a logical switch](#) with Switch Embedded Teaming (SET) enabled, so that the switch is virtualization aware. This switch is connected to the management logical network, and has all of the host virtual adapters that are required to provide access to the management network, or configure storage networking.

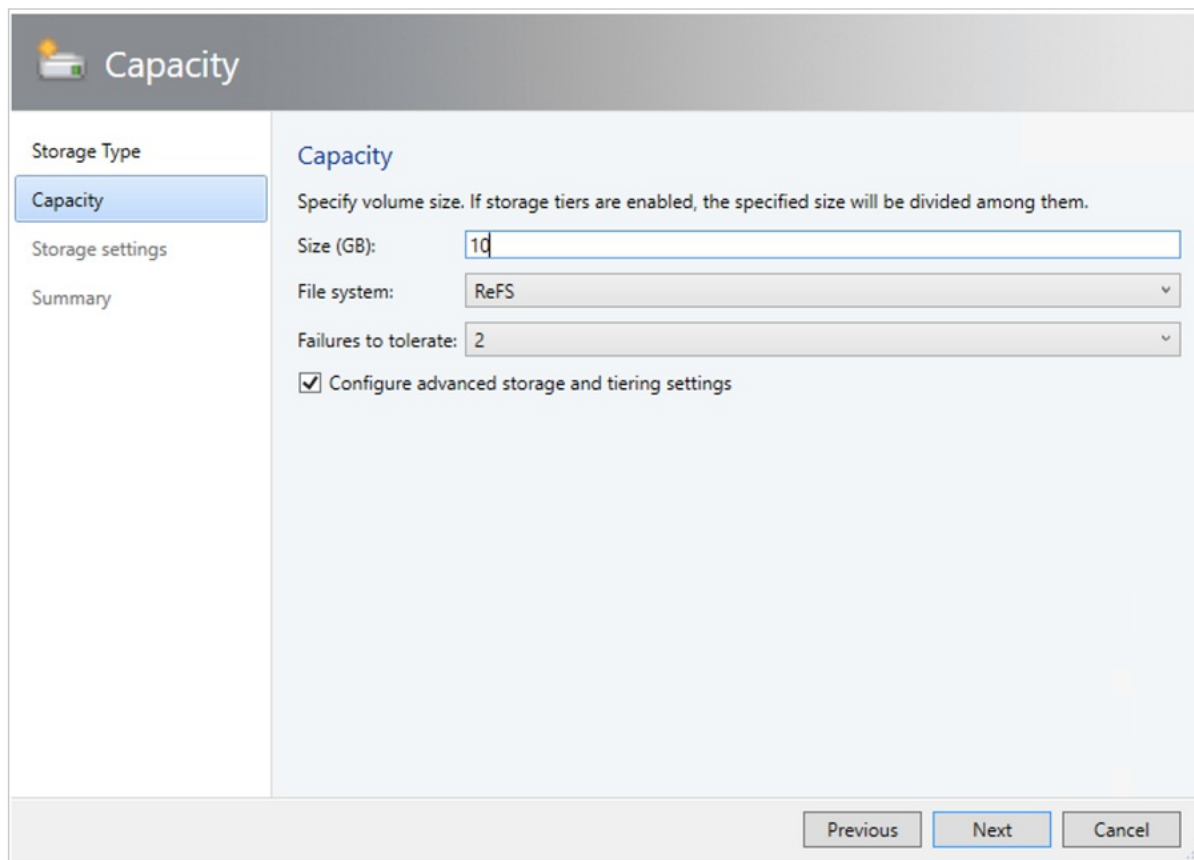
S2D relies on a network to communicate between hosts. RDMA-capable adapters are recommended.

3. [Create VM networks](#).

Step 3: Manage the pool and create CSVs

You can now modify the storage pool settings, and create virtual disks and CSVs.

1. Click **Fabric** > **Storage** > **Arrays**.
2. Right-click the cluster > **Manage Pool**, and select the storage pool that was created by default. You can change the default name, and add a classification.
3. To create a CSV, right-click the cluster > **Properties** > **Shared Volumes**.
4. In the Create Volume Wizard > **Storage Type**, specify the volume name, and select the storage pool.
5. In **Capacity**, you can specify the volume size, file system, and resiliency settings.



The screenshot shows the 'Capacity' step of the 'Create Volume Wizard'. The window has a title bar with a disk icon and the word 'Capacity'. On the left is a sidebar with three options: 'Storage Type' (selected), 'Storage settings', and 'Summary'. The main area is titled 'Capacity' and contains the following fields and options:

- A descriptive text: 'Specify volume size. If storage tiers are enabled, the specified size will be divided among them.'
- 'Size (GB):' with a text input field containing '10'.
- 'File system:' with a dropdown menu showing 'ReFS'.
- 'Failures to tolerate:' with a dropdown menu showing '2'.
- A checkbox labeled 'Configure advanced storage and tiering settings' which is checked.

At the bottom right, there are three buttons: 'Previous', 'Next' (highlighted with a blue border), and 'Cancel'.

6. Click **Configure advanced storage and tiering settings** to set up these options.

Create Volume Wizard

Storage settings

Storage Type
Capacity
Storage settings
Summary

Storage settings

Specify capacity and resiliency settings.

Storage tier count: 2

Storage tiers:

Media Type	Capacity (GB)	Resiliency
HDD	1	Mirror
HDD	9	Parity

Total capacity (GB): 10

File system allocation size

Allocation unit size (KB): 64

Previous Next Cancel

7. In **Summary**, verify settings and finish the wizard. A virtual disk will be created automatically when you create the volume.

If you use Powershell, the pool and the storage tier is automatically created with the "Enable-ClusterS2D autoconfig=true" option.

Step 4: Deploy VMs on the cluster

In a hyper-converged topology VMs can be directly deployed on the cluster. Their virtual hard disks are placed on the volumes you created using S2D. You [create and deploy these VMs](#) just as you would any other VM.

Next steps

- [Provision VMs](#)
- [Manage the cluster](#)

Deploy a Storage Spaces Direct disaggregated cluster in VMM

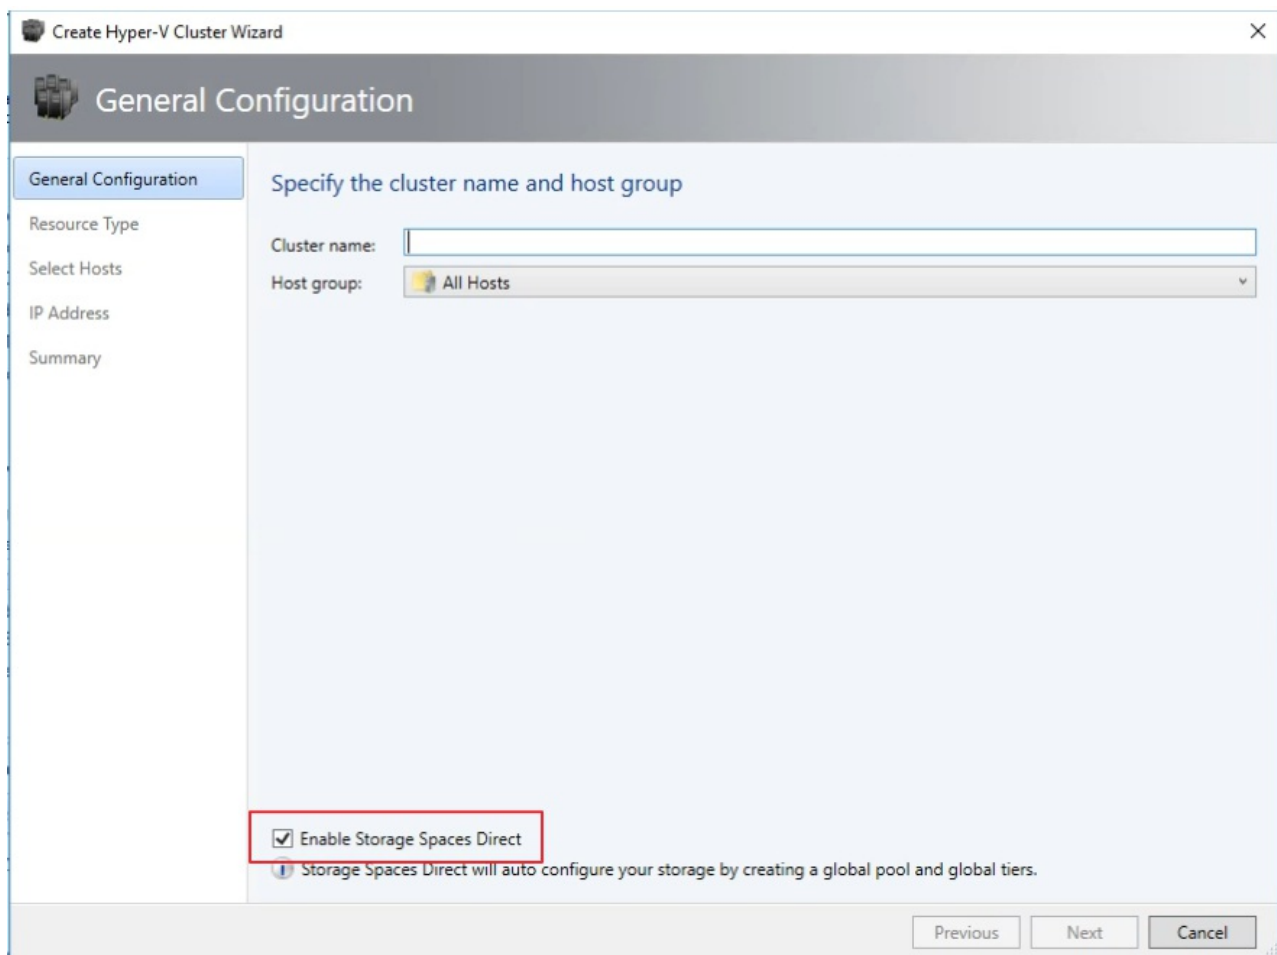
3/14/2019 • 2 minutes to read

Read this article to set up a disaggregated cluster running Storage Spaces Direct (S2D) in System Center - Virtual Machine Manager (VMM). [Learn more](#) about S2D.

You can deploy a disaggregated S2D cluster by provisioning a cluster running Hyper-V hosts, and a separate storage cluster running scale-out file server (SOFS) with S2D.

NOTE

You must enable S2D before adding the storage provider to VMM. To enable S2D, go to **General Configuration > Specify the cluster name and host group** and select the **Enable Storage Spaces Direct** option, as shown below:



After you enable a disaggregated cluster with S2D, VMM does the following:

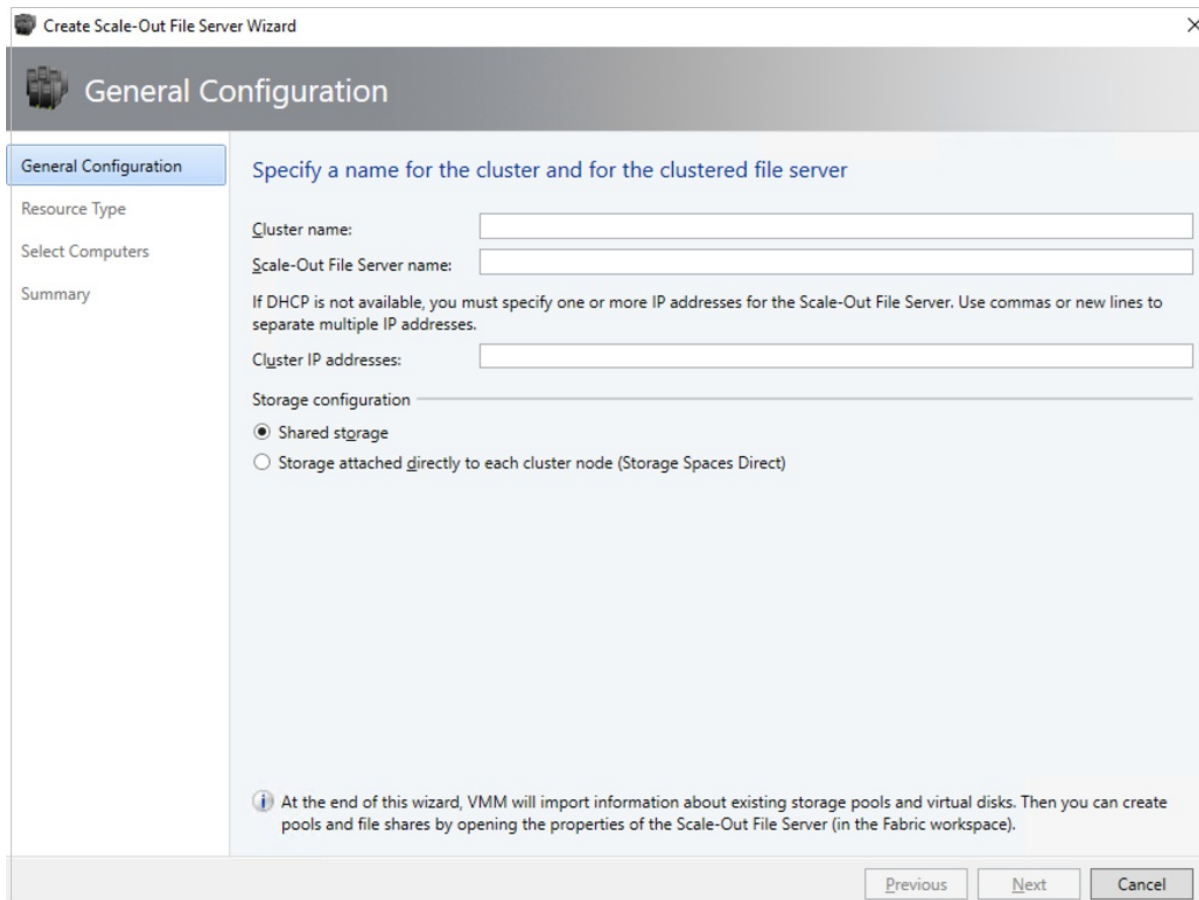
1. The File Server role and the Failover Clustering feature is enabled.
2. Storage replica and data deduplication is enabled.
3. The cluster is optionally validated and created.
4. S2D is enabled, and a storage array is created with the same name you provide in the wizard.

Step 1: Provision an SOFS cluster

You can provision an SOFS cluster from servers in the VMM fabric, or add an existing SOFS cluster to the fabric

Provision an SOFS cluster

1. Click **Fabric Resources > Create > File Server Cluster**.
2. In **General Configuration**, specify a cluster name, select a host group, and select **Storage attached directly to each cluster node (Storage Spaces Direct)**.



The screenshot shows the 'Create Scale-Out File Server Wizard' window, specifically the 'General Configuration' step. The window has a title bar with the text 'Create Scale-Out File Server Wizard' and a close button. Below the title bar is a header area with a server icon and the text 'General Configuration'. On the left side, there is a navigation pane with four items: 'General Configuration' (selected), 'Resource Type', 'Select Computers', and 'Summary'. The main area of the wizard is titled 'Specify a name for the cluster and for the clustered file server'. It contains several input fields: 'Cluster name:', 'Scale-Out File Server name:', and 'Cluster IP addresses:'. Below these fields is a section for 'Storage configuration' with two radio buttons: 'Shared storage' (selected) and 'Storage attached directly to each cluster node (Storage Spaces Direct)'. At the bottom of the main area, there is an information icon and a note: 'At the end of this wizard, VMM will import information about existing storage pools and virtual disks. Then you can create pools and file shares by opening the properties of the Scale-Out File Server (in the Fabric workspace)'. At the bottom right of the window, there are three buttons: 'Previous', 'Next', and 'Cancel'.

3. In **Resource Type**, specify the RunAs account with local admin permissions on the servers you want to add to the cluster, and specify whether to add existing Windows servers or bare-metal machines.
4. In **Cluster Nodes**, define a list of computers to add to the cluster.
5. In **Summary**, confirm the settings and then click **Finish**.

If you want to add additional nodes to the SOFS cluster, VMM automatically discovers any disks associated with the node. When you modify a storage pool and select the new disks to add, VMM makes those disks available to the hosts and VMs that use the share supported by that pool. [Learn more](#) about adding nodes to an SOFS.

Add an existing SOFS cluster with S2D enabled

1. Click **Fabric > Add Resources > Storage Devices**.
2. In the Add Resource Wizard, select **Windows-based File Server**.
3. In **Discovery Scope**, specify the cluster IP address or FQDN. Provide a Run As account with cluster access, and specify if the cluster is in a different domain.
4. In **Storage Device**, select the SOFS to add to the VMM fabric. You should only assign a classification to the pool after you add the provider.
5. In **Summary**, check the settings and complete the wizard.

Step 2: Manage the pool and create file shares

After the SOFS cluster is provisioned and managed by VMM, you can modify the storage pool, and create the

storage.

1. Click **Fabric** > **Storage** > **Arrays**.
2. Right-click the cluster > **Manage Pool**, and select the storage pool that was created by default. You can change the default name, and add a classification
3. After the pool appears with the new name if needed, click **Create File Share**.
4. In the Create File Share Wizard > **General**, specify a name for the share, and select the pool from which storage should be taken.
5. In **Capacity**, specify the share size and settings.
6. In **Summary**, verify the settings. After the share is created, a new CSV is added under the storage pool.

Step 3: Allocate the storage in Hyper-V

1. In the Hyper-V host properties > **Storage**, specify the file share path.
2. Now, you can create VMs that use this file share.

Next steps

- [Provision VMs](#)
- [Manage the cluster](#)

Manage Storage Spaces Direct clusters

3/19/2019 • 2 minutes to read

This article explains how to manage Storage Spaces Direct (S2D) clusters in the System Center - Virtual Machine Manager (VMM) fabric.

Configure cluster settings

You can [view and configure cluster settings](#), including cluster status, failure management, and available storage.

Add a node to a hyper-converged cluster

- You can [add a new node](#) on a hyper-converged S2D cluster in the VMM fabric. The new node can be an existing Hyper-V server, or a bare-metal physical server.

NOTE

Typically, S2D node requires RDMA, QoS and SET settings. To configure these settings for a node using bare metal computers, you can use the post deployment script capability in PCP. Here is the [sample PCP post deployment script](#). You can also use this script to configure RDMA, QoS and SET while adding a new node to an existing S2D deployment from bare metal computers.

- When you add a new node on a hyper-converged cluster, VMM automatically discovers disks on the new node, and enables S2D.
- VMM disables maintenance mode on disks, before adding them.

Control storage resources with QoS

Quality-of-service (QoS) in [Windows Server](#) provides a way to specify minimum and maximum resources that can be assigned to Hyper-V VMs using SOFS storage. QoS mitigates "noisy neighbor" issues, and ensures that a single VM doesn't consume all storage resources.

[Set up QoS policies](#) for a file server, or for specific virtual disks on the server.

7/9/2019 • 2 minutes to read

Assign storage QoS policy for clusters

VMM 1801 and later supports QoS on all managed clusters and also SOFS, running on Windows Server 2016 and beyond.

1. Click **Fabric > Storage > QoS Policies > Create Storage QoS Policy**.
2. In the wizard > **General**, specify a policy name.
3. In **Policy Settings**, specify how the policy should apply. Select **All virtual disk instances share resources** to specify that the policy should be applied to all virtual disks on the file server (pooled, single instance). Select **Resources allocated to each virtual disk instance** to specify that the policy is applied separately to each specified virtual disk (multi instance). Specify the minimum and maximum IOPS. A setting of 0 means that no policy is enforced.
4. In **Scope**, select the managed cluster under **Clusters** to which you want to apply the policy.

After upgrade, existing deployments which are managing their QoS with VMM, can seamlessly migrate to the new

QoS scoping based on the cluster name.

PowerShell cmdlets

The following new parameters are added:

AFFECTED CMDLET	PARAMETER	DETAILS
New-SCStorageQoSPolicy	-HostCluster	Specifies an array of HostCluster objects for QoS policy scope. Optional.
New-SCVirtualDiskDrive	-StorageQoSPolicy	Allows you to select the storage QoS policy for the Virtual Disk Drive. Optional.
Set-SCStorageQoSPolicy	-HostCluster	Specifies an array of HostCluster objects to be added in the QoS policy scope. Optional.
Get-SCStorageQoSPolicy	-HostCluster	Specifies a HostCluster object for which we want to query the QoS policies. Optional.

Assign a storage QoS Policy from templates

Templates usage is a common way for deploying VMs and Services on a cloud.

With VMM 1801 and later, you can select storage QoS policies from a template as well. For information on how to assign storage QoS policies from templates, see the related procedure in [create a VM template](#) article.

Next steps

[Manage QoS](#)

System requirements for System Center Virtual Machine Manager

3/14/2019 • 14 minutes to read

This article details the system requirements for System Center 2019 - Virtual Machine Manager (VMM).

VMM 2019 system requirements

The following sections describe the scalability information, hardware, software, and SQL Server requirements for VMM 2019, and summarize the support for the servers managed in the VMM fabric.

Capacity limits

The following table provides the scale limits that were tested for System Center Virtual Machine Manager 2019. There are various factors that impact the scale limits, such as hardware configuration, network, topology, and others.

The [planning guide](#) provides the details about how these factors can be adapted to specific requirements.

ENTITY	RECOMMENDED MAXIMUM COUNT
Physical hosts	1000
Virtual Machines	25000
Services	1000
User roles	1000
Clouds	20
Virtual networks	2000
logical networks	20
Library resources	1000
Library Objects (templates, profiles)	100

Hardware

HARDWARE	VMM SERVER	VMM DATABASE	VMM LIBRARY	VMM CONSOLE
Processor (minimum)	8 core Pentium 4, 2 GHz (x64)	8 core Pentium 4, 2.8 GHz	4 core Pentium 4, 2.8 GHz	2 core Pentium 4, 1 GHz CPU
Processor (recommended)	16-core, 2.66 GHz CPU	16 core 2.6 GHz CPU	4 core 2.8 GHz CPU	2 core 2 GHz CPU

HARDWARE	VMM SERVER	VMM DATABASE	VMM LIBRARY	VMM CONSOLE
RAM (minimum)	4 GB	8 GB	2 GB	4 GB
RAM (recommended)	16 GB	16 GB	4 GB	4 GB
Hard drive (minimum)	4 GB	50 GB	Based on size/amount of stored files	10 GB
Hard drive (recommended)	10 GB	200 GB	Based on size/amount of stored files	10 GB

Server operating system

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2016 (Server Core)	Y	N	If supported by SQL Server version
Windows Server 2016 (with desktop experience)	Y	Y	If supported by SQL Server version
Windows Server 2019 (with desktop experience)	Y	Y	If supported by SQL Server version
Windows Server 2019 (Server Core)	Y	Y	If supported by SQL Server version

NOTE

Ensure that VMM server operating system is same as the managed host operating system in case of deployment of Hyper Converged Infrastructure.

VMM console operating system

OPERATING SYSTEM	SUPPORTED
Windows 10 Enterprise	Y
Window Server 2016 Standard, Datacenter	Y
Windows Server 2019 (with desktop experience)	Y
Window Server 2019 Standard, Datacenter, Server Core with FOD	Y

SQL Server

NOTE

- For the supported versions of SQL, use the service packs that are currently in support by Microsoft.
- For the below supported SQL versions, Standard, Enterprise and Datacenter (64-bit) editions are supported, based on the availability.

SQL VERSION	SUPPORTED
SQL Server 2016 and SPs as detailed here	Y
SQL Server 2017 as detailed here	Y
SQL Server command line utilities	<p>Install the SQL Server 2016 Command-Line Utilities from the Microsoft® SQL Server® 2016 Feature Pack or</p> <p>Install the SQL Server 2017 Command-Line Utilities from the Microsoft® SQL Server® 2017 Feature Pack.</p> <p>Not required for VMM installation..</p>

Virtualization

VM	SUPPORTED
VMM management server	<p>The VMM management server can be installed on a VM.</p> <p>If you use dynamic memory, set the start RAM of the VM to at least 2048 MB.</p> <p>Don't install on a server running Hyper-V.</p> <p>You can deploy the VMM management server (physical or VM) in a highly available cluster.</p>
VMM console	You can install the VMM console on a VM.

Installation components

These components should be installed on the server, before you install VMM.

COMPONENT	VMM SERVER	VMM CONSOLE
Active Directory	<p>The VMM management server must be a domain member.</p> <p>The computer name should not exceed 15 characters.</p>	A computer with the VMM console installed should be a domain member.
Windows ADK	Download Windows ADK for Windows 10 and download windows PE Add on for ADK	Not applicable
PowerShell	PowerShell 5.0	PowerShell 4.0, 5.0

COMPONENT	VMM SERVER	VMM CONSOLE
.NET (minimum)	4.6	4.5

Servers in the VMM fabric

OPERATING SYSTEM	HYPER-V HOST	SOFS	UPDATE SERVER	PXE SERVER
Windows Server 2016	Y	Y	N	N
Windows Server 2016 (with desktop experience)	Y	Y	Y	Y
Windows Server 2019 (Data Center)	Y	Y	Y	Y
Windows Server 2019 (Server Core)	Y	Y	N	N

VMWare servers in the VMM fabric

VMWARE	SUPPORTED
ESX	ESX/ESXi 5.1, 5.5, 6.0, 6.5
vCenter	5.1, 5.5, 5.8, 6.0, 6.5
Supported	Features and limitations

VMs in the VMM fabric

GUEST OPERATING SYSTEM	SUPPORTED
Hyper-V VMs	Any guest running on supported Hyper-V hosts. Learn more about support for 2019 and earlier versions .
VMWare VM	Any VM running on supported VMware servers. Learn more

This article details the system requirements for System Center 1807 - Virtual Machine Manager (VMM).

VMM 1807 system requirements

The following sections describe the scalability information, hardware, software, and SQL Server requirements for VMM 1807, and summarize the support for the servers managed in the VMM fabric.

Capacity limits

The following table provides the scale limits that were tested for System Center Virtual Machine Manager. There are various factors that impact the scale limits, such as hardware configuration, network, topology, and others.

The [planning guide](#) provides the details about how these factors can be adapted to specific requirements.

ENTITY	RECOMMENDED MAXIMUM COUNT
Physical hosts	1000
Virtual Machines	25000
Services	1000
User roles	1000
Clouds	20
Virtual networks	2000
logical networks	20
Library resources	1000
Library Objects (templates, profiles)	100

Hardware

HARDWARE	VMM SERVER	VMM DATABASE	VMM LIBRARY	VMM CONSOLE
Processor (minimum)	8 core Pentium 4, 2 GHz (x64)	8 core Pentium 4, 2.8 GHz	4 core Pentium 4, 2.8 GHz	2 core Pentium 4, 1 GHz CPU
Processor (recommended)	16-core, 2.66 GHz CPU	16 core 2.6 GHz CPU	4 core 2.8 GHz CPU	2 core 2 GHz CPU
RAM (minimum)	4 GB	8 GB	2 GB	4 GB
RAM (recommended)	16 GB	16 GB	4 GB	4 GB
Hard drive (minimum)	4 GB	50 GB	Based on size/amount of stored files	10 GB
Hard drive (recommended)	10 GB	200 GB	Based on size/amount of stored files	10 GB

Server operating system

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 R2 Standard/Datacenter (Server Core)	N	Y	If supported by SQL Server version

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 R2 Standard/Datacenter (with desktop experience)	N	Y	If supported by SQL Server version
Windows Server 2016 (Server Core)	Y	N	If supported by SQL Server version
Windows Server 2016 (with desktop experience)	Y	Y	If supported by SQL Server version
Windows Server 1709	Y	Y	Y

VMM console operating system

OPERATING SYSTEM	SUPPORTED
Windows 10 Enterprise	Y
Windows Server 2012 R2 Standard, Datacenter	Y
Window Server 2016 Standard, Datacenter	Y

SQL Server

NOTE

- For the supported versions of SQL, use the service packs that are currently in support by Microsoft.
- For the below supported SQL versions, Standard, Enterprise and Datacenter (64-bit) editions are supported, based on the availability.

SQL VERSION	SUPPORTED
SQL Server 2008	N
SQL Server 2012 and SPs as detailed here	Y
SQL Server 2014 and SPs as detailed here	Y
SQL Server 2016 and SPs as detailed here	Y
SQL Server 2017 as detailed here	Y

SQL VERSION	SUPPORTED
SQL Server command line utilities	<p>Install on VMM server if you want to deploy SQL Server data-tier apps in the VMM fabric.</p> <p>Install SQL Server 2012 Command-Line Utilities from the Microsoft® SQL Server® 2012 Feature Pack or Install the SQL Server 2014 Command-Line Utilities from the Microsoft® SQL Server® 2014 Feature Pack or Install the SQL Server 2016 Command-Line Utilities from the Microsoft® SQL Server® 2016 Feature Pack</p> <p>Not required for VMM installation.</p>

Virtualization

VM	SUPPORTED
VMM management server	<p>The VMM management server can be installed on a VM.</p> <p>If you use dynamic memory, set the start RAM of the VM to at least 2048 MB.</p> <p>Don't install on a server running Hyper-V.</p> <p>You can deploy the VMM management server (physical or VM) in a highly available cluster.</p>
VMM console	You can install the VMM console on a VM.

Installation components

These components should be installed on the server, before you install VMM.

COMPONENT	VMM SERVER	VMM CONSOLE
Active Directory	<p>The VMM management server must be a domain member.</p> <p>The computer name should not exceed 15 characters.</p>	A computer with the VMM console installed should be a domain member.
Windows ADK	Download Windows ADK for Windows 10	Not applicable
PowerShell	PowerShell 5.0	PowerShell 4.0, 5.0
.NET	4.6	4.5, 4.5.1, 4.5.2, 4.6

Servers in the VMM fabric

OPERATING SYSTEM	HYPER-V HOST	SOFS	UPDATE SERVER	PXE SERVER
Windows Server 2012 Standard/Datacenter	N	N	N	N
Windows Server 2012 R2 Standard/Datacenter	Y	Y	Y	Y
Windows Server 2016	Y	Y	N	N
Windows Server 2016 (with desktop experience)	Y	Y	Y	Y
Windows Server 1709	Y	Y	N	N

VMWare servers in the VMM fabric

VMWARE	SUPPORTED
ESX	ESX/ESXi 5.1, 5.5, 6.0, 6.5
vCenter	5.1, 5.5, 5.8, 6.0, 6.5
Supported	Features and limitations

VMs in the VMM fabric

GUEST OPERATING SYSTEM	SUPPORTED
Hyper-V VMs	Any guest running on supported Hyper-V hosts. Learn more about support for 2016 and earlier versions .
VMWare VM	Any VM running on supported VMware servers. Learn more

This article details the system requirements for System Center 1801 - Virtual Machine Manager (VMM).

This article details the system requirements for System Center 2016 - Virtual Machine Manager (VMM).

VMM 1801 system requirements

The following sections describe the scalability information, hardware, software, and SQL Server requirements for VMM 1801, and summarize the support for the servers managed in the VMM fabric.

Capacity limits

The following table provides the scale limits that were tested for System Center 1801 Virtual Machine Manager. There are various factors that impact the scale limits, such as hardware configuration, network, topology, and others.

The [planning guide](#) provides the details about how these factors can be adapted to specific requirements.

ENTITY	RECOMMENDED MAXIMUM COUNT
Physical hosts	1000
Virtual Machines	25000
Services	1000
User roles	1000
Clouds	20
Virtual networks	2000
logical networks	20
Library resources	1000
Library Objects (templates, profiles)	100

Hardware

HARDWARE	VMM SERVER	VMM DATABASE	VMM LIBRARY	VMM CONSOLE
Processor (minimum)	8 core Pentium 4, 2 GHz (x64)	8 core Pentium 4, 2.8 GHz	4 core Pentium 4, 2.8 GHz	2 core Pentium 4, 1 GHz CPU
Processor (recommended)	16-core, 2.66 GHz CPU	16 core 2.6 GHz CPU	4 core 2.8 GHz CPU	2 core 2 GHz CPU
RAM (minimum)	4 GB	8 GB	2 GB	4 GB
RAM (recommended)	16 GB	16 GB	4 GB	4 GB
Hard drive (minimum)	4 GB	50 GB	Based on size/amount of stored files	10 GB
Hard drive (recommended)	10 GB	200 GB	Based on size/amount of stored files	10 GB

Server operating system

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 R2 Standard/Datacenter (Server Core)	N	Y	If supported by SQL Server version

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 R2 Standard/Datacenter (with desktop experience)	N	Y	If supported by SQL Server version
Windows Server 2016 (Server Core)	Y	N	If supported by SQL Server version
Windows Server 2016 (with desktop experience)	Y	Y	If supported by SQL Server version
Windows Server 1709	Y	Y	Y

VMM console operating system

OPERATING SYSTEM	SUPPORTED
Windows 10 Enterprise	Y
Windows Server 2012 R2 Standard, Datacenter	Y
Window Server 2016 Standard, Datacenter	Y

SQL Server

NOTE

- For the supported versions of SQL, use the service packs that are currently in support by Microsoft.
- For the below supported SQL versions, Standard, Enterprise and Datacenter (64-bit) editions are supported, based on the availability.

SQL VERSION	SUPPORTED
SQL Server 2008	N
SQL Server 2012 and SPs as detailed here	Y
SQL Server 2014 and SPs as detailed here	Y
SQL Server 2016 and SPs as detailed here	Y

SQL VERSION	SUPPORTED
SQL Server command line utilities	<p>Install on VMM server if you want to deploy SQL Server data-tier apps in the VMM fabric.</p> <p>Install SQL Server 2012 Command-Line Utilities from the Microsoft® SQL Server® 2012 Feature Pack or Install the SQL Server 2014 Command-Line Utilities from the Microsoft® SQL Server® 2014 Feature Pack or Install the SQL Server 2016 Command-Line Utilities from the Microsoft® SQL Server® 2016 Feature Pack</p> <p>Not required for VMM installation.</p>

Virtualization

VM	SUPPORTED
VMM management server	<p>The VMM management server can be installed on a VM.</p> <p>If you use dynamic memory, set the start RAM of the VM to at least 2048 MB.</p> <p>Don't install on a server running Hyper-V.</p> <p>You can deploy the VMM management server (physical or VM) in a highly available cluster.</p>
VMM console	You can install the VMM console on a VM.

Installation components

These components should be installed on the server, before you install VMM.

COMPONENT	VMM SERVER	VMM CONSOLE
Active Directory	<p>The VMM management server must be a domain member.</p> <p>The computer name should not exceed 15 characters.</p>	A computer with the VMM console installed should be a domain member.
Windows ADK	Download Windows ADK for Windows 10	Not applicable
PowerShell	PowerShell 5.0	PowerShell 4.0, 5.0
.NET	4.6	4.5, 4.5.1, 4.5.2, 4.6

Servers in the VMM fabric

OPERATING SYSTEM	HYPER-V HOST	SOFS	UPDATE SERVER	PXE SERVER
Windows Server 2012 Standard/Datacenter	N	N	N	N
Windows Server 2012 R2 Standard/Datacenter	Y	Y	Y	Y
Windows Server 2016	Y	Y	N	N
Windows Server 2016 (with desktop experience)	Y	Y	Y	Y
Windows Server 1709	Y	Y	N	N

VMWare servers in the VMM fabric

VMWARE	SUPPORTED
ESX	ESX/ESXi 5.1, 5.5, 6.0
vCenter	5.1, 5.5, 5.8, 6.0
Supported	Features and limitations

VMs in the VMM fabric

GUEST OPERATING SYSTEM	SUPPORTED
Hyper-V VMs	Any guest running on supported Hyper-V hosts. Learn more about support for 2016 and earlier versions .
VMWare VM	Any VM running on supported VMware servers. Learn more

VMM 2016 system requirements

The following sections describe the scalability information, hardware, software, and SQL Server requirements for VMM 2016, and summarize the support for the servers managed in the VMM fabric.

Capacity limits

The following table provides the scale limits that were tested for System Center 2016 Virtual Machine Manager. There are various factors that impact the scale limits, such as hardware configuration, network, topology, and others.

The [planning guide](#) provides the details about how these factors can be adapted to specific requirements.

ENTITY	RECOMMENDED MAXIMUM COUNT
Physical hosts	1000
Virtual Machines	25000
Services	1000
User roles	1000
Clouds	20
Virtual networks	2000
logical networks	20
Library resources	1000
Library Objects (templates, profiles)	100

Hardware

HARDWARE	VMM SERVER	VMM DATABASE	VMM LIBRARY	VMM CONSOLE
Processor (minimum)	8 core Pentium 4, 2 GHz (x64)	8 core Pentium 4, 2.8 GHz	4 core Pentium 4, 2.8 GHz	2 core Pentium 4, 1 GHz CPU
Processor (recommended)	16-core, 2.66 GHz CPU	16 core 2.6 GHz CPU	4 core 2.8 GHz CPU	2 core 2 GHz CPU
RAM (minimum)	4 GB	8 GB	2 GB	4 GB
RAM (recommended)	16 GB	16 GB	4 GB	4 GB
Hard drive (minimum)	4 GB	50 GB	Based on size/amount of stored files	10 GB
Hard drive (recommended)	10 GB	200 GB	Based on size/amount of stored files	10 GB

Server operating system

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 Standard/Datacenter	N	N	If supported by SQL Server version
Windows Server 2012 R2 Standard/Datacenter (Server Core)	N	Y	If supported by SQL Server version

OPERATING SYSTEM	VMM SERVER	REMOTE VMM LIBRARY	REMOTE VMM DATABASE
Windows Server 2012 R2 Standard/Datacenter (with desktop experience)	N	Y	If supported by SQL Server version
Windows Server 2016 (Server Core)	Y	N	If supported by SQL Server version
Windows Server 2016 (with desktop experience)	Y	Y	If supported by SQL Server version
Windows Server 2016 Nano	N	N	If supported by SQL Server version

VMM console operating system

OPERATING SYSTEM	SUPPORTED
Windows 7	N
Windows 8	N
Windows 8.1	Y
Windows 10 Enterprise	Y
Windows Server 2008 R2 with SP1 onwards	N
Windows Server 2012	Y
Windows Server 2012 R2 Standard, Datacenter	Y
Window Server 2016 Standard, Datacenter	Y

SQL Server

NOTE

- For the supported versions of SQL, use the service packs that are currently in support by Microsoft.
- For the below supported SQL versions, Standard, Enterprise and Datacenter (64-bit) editions are supported, based on the availability.

SQL VERSION	SUPPORTED
SQL Server 2008	N
SQL Server 2012 and SPs as detailed here	Y
SQL Server 2014 and SPs as detailed here	Y

SQL VERSION	SUPPORTED
SQL Server 2016 and SPs as detailed here	Y
SQL Server command line utilities	<p>Install on VMM server if you want to deploy SQL Server data-tier apps in the VMM fabric.</p> <p>Install SQL Server 2012 Command-Line Utilities from the Microsoft® SQL Server® 2012 Feature Pack or Install the SQL Server 2014 Command-Line Utilities from the Microsoft® SQL Server® 2014 Feature Pack</p> <p>Not required for VMM installation.</p>

Virtualization

VM	SUPPORTED
VMM management server	<p>The VMM management server can be installed on a VM.</p> <p>If you use dynamic memory, set the start RAM of the VM to at least 2048 MB.</p> <p>Don't install on a server running Hyper-V.</p> <p>You can deploy the VMM management server (physical or VM) in a highly available cluster.</p>
VMM console	You can install the VMM console on a VM.

Installation components

These components should be installed on the server, before you install VMM.

COMPONENT	VMM SERVER	VMM CONSOLE
Active Directory	<p>The VMM management server must be a domain member.</p> <p>The computer name should not exceed 15 characters.</p>	A computer with the VMM console installed should be a domain member.
Windows ADK	Download Windows ADK for Windows 10	Not applicable
PowerShell	PowerShell 5.0	PowerShell 4.0, 5.0
.NET	4.6	4.5, 4.5.1, 4.5.2, 4.6

Servers in the VMM fabric

OPERATING SYSTEM	HYPER-V HOST	SOFS	UPDATE SERVER	PXE SERVER
Windows Server 2012 Standard/Datacenter	N	N	N	N
Windows Server 2012 R2 Standard/Datacenter	Y	Y	Y	Y
Windows Server 2016	Y	Y	N	N
Windows Server 2016 (with desktop experience)	Y	Y	Y	Y
Windows Server 2016 Nano	Y	Y	N	N

VMWare servers in the VMM fabric

VMWARE	SUPPORTED
ESX	ESX/ESXi 5.1, 5.5, 6.0 UR6 and later - ESX/ESXi 5.1, 5.5, 6.0, 6.5
vCenter	5.1, 5.5, 5.8, 6.0 UR6 and later - 5.1, 5.5, 5.8, 6.0, 6.6
Supported	Features and limitations

VMs in the VMM fabric

GUEST OPERATING SYSTEM	SUPPORTED
Hyper-V VMs	Any guest running on supported Hyper-V hosts. Learn more about support for 2016 and earlier versions .
VMWare VM	Any VM running on supported VMware servers. Learn more

Next steps

[Plan VMM installation](#)

Plan VMM installation

6/14/2019 • 11 minutes to read

This article helps you to plan all the elements required for a successful System Center - Virtual Machine Manager (VMM) installation and includes information for releases VMM 2016 and later. Please use these requirements as applicable for the VMM version you plan to install.

For additional information on the supported versions of hardware and software, see the system requirements article for the version you install.

Deployment requirements

Verify the following [system requirements](#):

- **VMM management server:** Verify hardware and operating system requirements.
- **SQL Server:** Review supported SQL Server versions
- **VMM console:** Review operating system requirements and if you want to run the VMM console on a separate computer.
- **VMM library:** Review the hardware requirements for remote VMM library shares.
- **Virtualization hosts:** Review the supported operating systems for Hyper-V and SOFS servers in the VMM fabric. Review requirements for VMware servers.
- **Other fabric servers:** Review the supported operating systems for update and PXE (used for bare metal deployment) servers.

Additional deployment requirements

COMPONENT	DETAILS
Command line utilities for SQL Server	SQL Server 2014 feature pack for release earlier to 2019 , 2016/2017 feature pack for 2019 If you want to deploy VMM services using SQL Server data-tier apps, install the related command-line utilities on the VMM management server. The version you install should match the SQL Server version. You don't have to install these to install VMM.
Windows Assessment and Deployment Kit (ADK)	Windows ADK for Windows 10. You can install from setup, or download it . You only need the Deployment Tools and Windows Preinstallation Environment options.
Guest operating system	Windows operating systems supported by Hyper-V . Linux (CentOS, RHEL, Debian, Oracle Linux, SUSE, Ubuntu)
PowerShell	Supported versions
.NET	Supported versions

COMPONENT	DETAILS
Host agent	VMM 2016/1801/1807/2019 Needed for hosts managed in VMM.
Monitoring	System Center Operations Manager 2016/1801. You also need SQL Server Analysis Services 2014 or a later version.
VMware	vCenter 5.1, 5.5, 5.8, 6.0, 6.5 ESX 5.5, ESX 6.0, ESX 6.5 vCenter and ESX servers running these versions can be managed in VMM.
Bare metal provisioning	System Management Architecture for Server Hardware (SMASH) (v1 or higher) over WS-MAN. Intelligent Platform Interface 1.5 or higher Data Center Manager Interface (DCMI) 1.0 or higher. Required to discover and deploy physical bare metal servers.

SPN

If the VMM user installing VMM, or running VMM setup, doesn't have permissions to write the service principal name (SPN) for the VMM server in Active Directory, setup will finish with a warning. If the SPN isn't registered, other computers running the VMM console won't be able to connect to the management server, and you won't be able to deploy a Hyper-V host on a bare metal computer in the VMM fabric. To avoid this issue, you need to register the SPN as a domain administrator before you install VMM, as follows:

1. Run these commands from <SystemDrive>\Windows\System32>, as a domain administrator:

```
setspn -s -u SCVMM/<MachineBIOSName> <VMMSserviceAccount> setspn -s -u SCVMM/<MachineFQDN>
<VMMSserviceAccount>
```

For a cluster, <MachineBIOSName> should be <ClusterBIOSName> and <MachineFQDN> should be <ClusterFQDN>

2. On the VMM server (or on each node in a cluster), in the registry, navigate to **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine Manager Server\Setup**.
3. Set **VmmServicePrincipalNames** to **SCVMM/<MachineBIOSName>,SCVMM/<MachineFQDN>**.
For a cluster: **SCVMM/<ClusterBIOSName>,SCVMM/<ClusterFQDN>**.

If you can't do this, you can also register the SPN during VMM installation. A domain administrator can provide the SPN write permissions to VMM service user or setup user. Note that this approach isn't the preferred one. The permission allows the delegated user to register any servicePrincipalName, with no restrictions. Hence, the delegated user should be highly trusted, and the account credentials must be kept secure. To do this:

1. Run **adsiedit** as a domain administrator.
2. Navigate to find the VMM service user. Right-click **Properties** > **Security** > **Advanced**. Then click **Add**, and in **Select a principal**, specify user who will be granted the permissions.
3. Select **Write servicePrincipalName** > **OK**.

When you install VMM with this user account SPN will be registered.

VMM management server

- You can't run the VMM management server on Nano server (applicable to releases prior to 2019).
- The management server computer name cannot exceed 15 characters.
- Don't install the VMM management server, or other System Center components other than agents, on servers running Hyper-V.
- You can install the VMM management server on a VM. If you do, and you use the Dynamic Memory feature of Hyper-V, then you must set the startup RAM for the virtual machine to be at least 2,048 megabytes (MB).
- If you want to manage more than 150 hosts, we recommend that you use a dedicated computer for the VMM management server and do the following:
 - Add one or more remote computers as library servers, and do not use the default library share on the VMM management server.
 - Don't run the SQL Server instance on the VMM management server.
- For high availability, the VMM management server can be installed on a failover cluster. [Learn more.](#)

SQL Server and database

- The instance of SQL Server that you are using must allow for case-insensitive database objects.
- The SQL Server's computer name cannot exceed 15 characters in length.
- If the VMM management server and the SQL Server computer are not members of the same Active Directory domain, then a two-way trust must exist between the two domains.
- When you install SQL Server, select the **Database Engine Services** and **Management Tools - Complete** features.
- You can perform an in-place upgrade to a supported version of SQL Server (without moving the VMM database). Make sure no jobs are running when you perform the upgrade, or jobs may fail and may need to be restarted manually.
- For the VMM database, for better performance, do not store database files on the disk that is used for the operating system.
- If you are using Software Defined Networking (SDN) in VMM, then all networking information is stored in the VMM database. Because of this, you might want to consider high availability for the VMM database, using the following guidelines:
 - Failover clustering is supported and is the recommended configuration for availability within a single geographical area or datacenter. [Read more.](#)
 - Use of Always On Availability Groups in Microsoft SQL Server is supported, but it's important to review the differences between the two availability modes, synchronous-commit and asynchronous-commit. [Learn more.](#)
 - With asynchronous-commit mode, the replica of the database can be out of date for a period of time after each commit. This can make it appear as if the database were back in time which might cause loss of customer data, inadvertent disclosure of information, or possibly elevation of privilege.
 - You can use synchronous-commit mode as a configuration for remote-site availability scenarios.
- The SQL Server service must use an account that has permission to access Active Directory Domain Services (AD DS). For example, you can specify the Local System Account, or a domain user account. Do not specify a local user account.
- You don't need to configure collation. During deployment, Setup automatically configures CI collation according to the language of the server operating system.

NOTE

Target database collate and the VMM server collate must be the same if the database is part of SQL Always On (AO) group. To check the current culture on your VMM machine, use the [Get-Culture](#) PowerShell command. Learn more about [VMM server collate mappings](#).

- Dynamic port is supported.
- If you want to create the VMM database prior to VMM installation:
 - Make sure you have permissions or create a SQL database, or ask the SQL Server admin to do it.
 - Configure the database as follows:
 1. Create a new database with settings: Name: VirtualManagerDB; Collation: Latin1_General_100_CI_AS, but aligned with the specific SQL Server instance collation.
 2. Grant db_owner permissions for the database to the VMM service account.
 3. In VMM setup you'll select the option to use an existing database and specify the database details and VMM service account as the database user.

Library server

- If you run the library server on the VMM management server, then you must provide additional hard disk space to store objects. The space required varies, based on the number and size of the objects you store.
- The library server is where VMM stores items such as virtual machine templates, virtual hard disks, virtual floppy disks, ISO images, scripts, and stored virtual machines. The optimal hardware requirements that are specified for a VMM library server vary, depending on the quantity and size of these files. You will need to check CPU usage, and other system state variables to determine what works best in your environment.
- If you want to manage Virtual hard disks in the .vhdx file format, the VMM library server must run Windows Server 2012 or later.
- VMM does not provide a method for replicating physical files in the VMM library or a method for transferring metadata for objects that are stored in the VMM database. Instead, if necessary, you need to replicate physical files outside of VMM, and you need to transfer metadata by using scripts or other means.
- VMM does not support file servers that are configured with the case-sensitive option for Windows Services for UNIX, because the Network File System (NFS) case control is set to **Ignore**.

Account and domain requirements

When you install VMM you need to configure the VMM service to use either the Local System account or a domain account or a Group Managed Service Account (gMSA).

Ensure the following before you prepare an account:

- It is not supported to change the identity of the Virtual Machine Manager service account after installation. This includes changing from the local system account to a domain account, from a domain account to the local system account, or changing the domain account to another domain account. To change the Virtual Machine Manager service account after installation, you must uninstall VMM (selecting the Retain data option if you want to keep the SQL Server database), and then reinstall VMM by using the new service account.
- If you specify a domain account, the account must be a member of the local Administrators group on the computer.
- If you specify a domain account, it is strongly recommended that you create an account that is specifically designated to be used for this purpose. When a host is removed from the VMM management server, the account that the System Center Virtual Machine Manager service is running under is removed from the local Administrators group of the host. If the same account is used for other purposes on the host, this can cause

unexpected results.

- If you plan to use shared ISO images with Hyper-V virtual machines, you must use a domain account.
- If you are using a disjointed namespace, you must use a domain account. For more information about disjointed namespaces, see [Naming conventions in Active Directory for computers, domains, sites, and OUs](#).
- If you are installing a highly available VMM management server, you must use a domain account.
- The computer on which you install the VMM management server must be a member of an Active Directory domain. In your environment you might have user accounts in one forest and your VMM servers and host in another. In this environment, you must establish a two-way trust between the two cross-forest domains. One-way trusts between cross-forest domains are not supported in VMM.
- To create and use gMSA, review the article on [gMSA](#) and create the gMSA as per the guidance available. Make sure that the servers on which the VMM Management service would be installed have permissions to retrieve the password of gMSA account.

NOTE

You do not need to specify the 'Service Principle Name (SPN)' when creating gMSA. VMM service sets the appropriate SPN for gMSA.

Distributed key management

By default, VMM encrypts some data in the VMM database by using the Data Protection Application Programming Interface (DPAPI). For example, Run As account credentials, passwords in guest operating system profiles, and product key information in virtual hard disks properties. Data encryption is tied to the specific computer on which VMM is installed, and the service account that VMM uses. If you move your VMM installation to another computer, VMM won't retain the encrypted data, and you'll need to enter it manually.

To ensure that VMM retains encrypted data across moves, you can use distributed key management to store encryption keys in Active Directory. If you move your VMM installation, VMM retains the encrypted data because the new VMM computer has access to the encryption keys in Active Directory. To set up distributed key management you should coordinate with your Active Directory administrator. Note that:

- You must create a container in AD DS before you install VMM. You can create the container by using ADSI Edit (installed from **Server Manager > Remote Server Administration Tools**).
- You create the container in the same domain as the user account with which you are installing VMM. If you specify that the VMM service uses a domain account, that account must be in the same domain. For example, if the installation account and the service account are both in the corp.contoso.com domain, you must create the container in that domain. So, if you want to create a container that is named VMMDKM, you specify the container location as CN=VMMDKM,DC=corp,DC=contoso,DC=com. The account with which you're installing VMM needs Full Control permissions to the container in AD DS. The permissions must apply to this object, and to all descendant objects.
- If you are installing a highly available VMM management server, you must use distributed key management to store encryption keys in Active Directory. You need distributed key management because if VMM fails over to a node, that node will need access to the encryption keys.
- When you configure the service account and distributed key in setup, you must type the location of the container in AD DS, for example: CN=VMMDKM,DC=corp,DC=contoso,DC=com

Next steps

[Install VMM](#)

Plan a highly available VMM deployment

12/7/2018 • 3 minutes to read

This article helps you to plan a highly available System Center - Virtual Machine Manager (VMM) deployment.

For resilience and scalability you can deploy VMM in high availability mode, as follows:

- Deploy the VMM management server in a failover cluster.
- Make library server file shares highly available
- Deploy the SQL Server VMM database as highly available.

Plan a highly available SQL Server deployment

- You should set up SQL Server before you deploy the VMM management servers.
- We recommend you use a highly available SQL Server installation on a failover cluster and configure SQL Server Always On availability groups. You shouldn't install SQL Server on the VMM cluster.
- Review the [best practices](#) for failover cluster node prerequisites.
- Always On availability groups are supported in VMM. Use **synchronous commit** for higher protection with more overhead. If you use **asynchronous-commit** mode the secondary database can lag behind the primary database making some data loss possible.
- The database server must be in the same domain as the VMM server, or in a domain with a two-way trust.
- Using a clustered database with VMM requires Kerberos authentication. To support this, the SQL Server instance must associate a Service Principal Name (SPN) with the account that SQL Server will be running on.

Plan a highly available VMM management server

- Don't install on a Hyper-V host parent partition. You can install VMM on a VM.
- Before you start you'll need to set up the VMM service account and distributed key management. [Learn more](#)
- Only one instance of VMM can be deployed to a failover cluster of up to 16 nodes.
- The user who creates the cluster has **Create Computers objects** permission to the OU or the container where the servers that will form the cluster reside. If this isn't possible ask a domain admin to pre-stage a cluster computer object for the cluster.
- Requirements for computers running as VMM management nodes:
 - All cluster nodes that will act as VMM servers must be running Windows Server 2016.
 - Each cluster node must be joined to a domain and must have a computer name that does not exceed 15 characters.
 - The VMM service network name must not exceed 15 characters.
 - Windows ADK needs to be installed on each computer. Install from setup or the [download center](#). Select **Deployment Tools** and **Windows Preinstallation Environment** when you install.
 - If you plan to deploy VMM services that use SQL Server data-tier applications, install the related command-line utilities on your VMM management server. The command line utility is available in the [SQL Server 2012 feature pack](#) or [SQL Server 2014 feature pack](#) or [SQL Server 2016 feature pack](#).

Plan a highly available VMM library

You can create highly-available library servers to ensure that file-based resources, templates, and profiles are resilient and available.

- VMM doesn't automatically create the VMM library as highly available when you deploy VMM in high availability mode. You need to create highly available library servers by deploying the library on a file server cluster.
- You'll need to set up a file server failover cluster. Deploying highly available library shares on the VMM cluster isn't supported.
- Computers you'll configure as file servers should be running Windows Server 2012 R2 or later. We recommend that all nodes have the same version of Windows.
- All nodes you want to add as file servers should be in the same domain.
- Make sure the hardware and software that you want to use for the library meets the system requirements.
- The user who creates the cluster has **Create Computers objects** permission to the OU or the contain where the servers that will form the cluster reside. If this isn't possible ask a domain admin to pre-stage a cluster computer object for the cluster.
- The account you use to create the cluster should be a domain user on all the computers you want to add as file server nodes.
- The library server can't be a scale-out file server (SOFS). It must be on a failover cluster that doesn't use the SOFS cluster role. This is because when you deploy the library the VMM agent is deployed on the host. For SOFS there are multiple hosts in a cluster provides shares which makes it complicated for agent deployment. When you have a standalone or clustered library server, you can leverage storage on SOFS by creating shares on it.
- You can deploy the library shares on a cluster with physical nodes, or a guest cluster.
- If you want to add clustered storage when you create the cluster make sure all computers can access the storage.
- If you want to deploy a distributed VMM library in different datacenters you'll need to set up a scheduled copy between the two library shares. No replication is available.

Next steps

- [Set up a highly available VMM deployment](#)

Identify VMM ports and protocols

5/17/2019 • 3 minutes to read

As part of your System Center Virtual Machine Manager (VMM) deployment, you need to allow access for the ports and protocols that the VMM server and components will use. It's important to plan this in advance. Some of the port settings are configured in VMM setup, and if you want to modify them after you set up VMM for the first time, you will need to reinstall to do so.

Set up exceptions

1. Identify where you need to create firewall exceptions, in accordance with the table below.
2. On the server you identify, click **Start > Next > Administrative Tools > Windows Firewall with Advanced Security**.
3. In the **Windows Firewall with Advanced Security on Local Computer** pane, click on **Inbound Rules**.
4. In **Actions**, click **New Rule**.
5. In the **New Inbound Rule Wizard > Rule Type**, click on **Port**, and then click on **Next**.
6. In **Protocol and Ports**, specify the port settings in accordance with the table below, and continue in the wizard to create the rule.

Port and protocol exceptions

CONNECT	PORT/PROTOCOL	DETAILS	CONFIGURE
VMM server to VMM agent on Windows Server-based hosts/remote library server	80: WinRM; 135: RPC; 139: NetBIOS; 445: SMB (over TCP)	Used by the VMM agent Inbound rule on hosts	Can't modify
VMM server to VMM agent on Windows Server-based hosts/remote library server	443:HTTPS	BITS data channel for file transfers Inbound rule on hosts	Modify in VMM setup
VMM server to VMM agent on Windows Server-based hosts/remote library server	5985:WinRM	Control channel Inbound rule on hosts	Modify in VMM setup
VMM server to VMM agent on Windows Server-based hosts/remote library server	5986:WinRM	Control channel (SSL) Inbound rule on hosts	Can't modify

CONNECT	PORT/PROTOCOL	DETAILS	CONFIGURE
VMM server to VMM guest agent (VM data channel)	443:HTTPS	<p>BITS data channel for file transfers</p> <p>Inbound rule on machines running the agent</p> <p>The VMM guest agent is a special version of the VMM agent. It's installed on VMs that are part of a service template, and on Linux VMs (with or without a service template).</p>	Can't modify
VMM server to VMM guest agent (VM control channel)	5985:WinRM	<p>Control channel</p> <p>Inbound rule on machines running the agent</p>	Can't modify
VMM host to host	443:HTTPS	<p>BITS data channel for file transfers</p> <p>Inbound rule on hosts and VMM server</p>	Modify in VMM setup
VMM server to VMware ESXi servers/Web Services	22:SFTP Inbound rule on hosts	Can't modify	
VMM server to load balancer	80:HTTP; 443:HTTPS	Channel used for load balancer management	Modify in load balancer provider
VMM server to remote SQL Server database	1433:TDS	<p>SQL Server listener</p> <p>Inbound rule on SQL Server</p>	Modify in VMM setup
VMM server to WSUS update servers	80/8530:HTTP; 443/8531:HTTPS	<p>Data and control channels</p> <p>Inbound rule on WSUS server</p>	Can't modify from VMM
VMM library server to Hyper-V hosts	443:HTTPS	<p>BITS data channel for file transfers</p> <p>Inbound rule on hosts - 443</p>	Modify in VMM setup
VMM console to VMM	WCF:8100 (HTTP); WCF:8101 (HTTPS); Net.TCP: 8102	Inbound rule on VMM console machine	Modify in VMM setup
VMM server to storage management service	WMI	Local call	
Storage management service to SMI-S provider	CIM-XML	Provider-specific	

CONNECT	PORT/PROTOCOL	DETAILS	CONFIGURE
VMM server to Baseboard Management Controller (BMC)	443: HTTP (SMASH over WS-Management)	Inbound rule on BMC device	Modify on BMC device
VMM server to Baseboard Management Controller (BMC)	623: IPMI	Inbound rule on BMC device	Modify on BMC device
VMM server to Windows PE agent	8101:WCF; 8103:WCF	8101 is used for control channel, 8103 is used for time sync	Modify in VMM setup
VMM server to WDS PXE provider		8102: WCF	Inbound rule on PXE server
VMM server to Hyper-V host in untrusted/perimeter domain	443:HTTPS (BITS)	BITS data channel for file transfers Inbound rule on VMM server	
Library server to Hyper-V host in untrusted/perimeter domain	443:HTTPS	BITS data channel for file transfers Inbound rule on VMM library	
VMM server to Windows file server	80: WinRM; 135: RPC; 139: NetBIOS; 445: SMB (over TCP)	Used by the VMM agent Inbound rule on file server	
VMM server to Windows file server	443:HTTPS	BITS used for file transfer Inbound rule on file server	
VMM server to Windows file server	5985/5986:WinRM	Control channel Inbound rule on file server	

NOTE

In addition to the above ports, VMM relies on the default dynamic port range for all its communication with Hyper-V hosts, file servers and library servers. [Learn more](#) about the dynamic port range. We recommend that you reconfigure the firewalls to allow traffic between servers in the dynamic port range of 49152 through 65535.

Next steps

You can modify some of these ports and protocols during [VMM installation](#).

Plan the VMM compute fabric

12/7/2018 • 5 minutes to read

This article describes how to plan the System Center - Virtual Machine Manager (VMM) compute fabric. The VMM compute fabric consists of the VMM library, virtualization hosts, host groups, and other infrastructure servers.

Plan the VMM library

Before you start:

- You should verify system requirements for the VMM library before you install VMM.
- VMM deploys the default library share on the VMM server. After setup is complete you can't remove or relocate the default library share so consider its location before you install VMM.
- If you're using a SAN the library server should have the same SAN as the hosts that use the library. This ensures that the library server and the hosts can access the same LUNs on the SAN for faster file transfers.
- If you connect to a library from virtualization hosts across a LAN the library server should be as close as possible to the hosts.
- If you're planning to add more library servers you can create library groups to organize them. You can use library groups to align servers with host groups in the VMM fabric. As a best practice you should align each library server with the host group that uses the resources of that library.

Plan virtualization hosts

VMM supports Hyper-V and VMware virtualization hosts. When you're adding, provisioning, and managing hosts in the VMM fabric consider these points:

- The topology of Hyper-V hosts. VMM can work with Hyper-V hosts that are located in the same domain as the VMM server, in a domain with a two-way trust, or in a domain without a two-way trust. VMM can also work with Hyper-V hosts that are in a perimeter network, or in a disjointed namespace.
- The topology of VMware hosts. VMM works with VMware hosts located anywhere in your environment.
- The number and type of guest operating systems running on the host.
- The system configuration of the VMs running on the host.
- The types of apps running on the guest operating systems.
- The VM workloads that will run on the host.
- The processor requirements for the host. You'll need enough processing capacity to run the VMs.
- The memory requirements for the host. After you use VMM to allocate host RAM to a VM that memory isn't available for other resources. You also need adequate memory to run the host operating system and any other apps.
- The storage requirements for the host. You need adequate storage for the host itself, as well as for the VMs running on it. Remember you'll need to take extra space into account for VM paging files, dynamically expanding virtual hard disks, saving the contents of VM RAM when the VM is in a saved state, and VM checkpoints.
- The network requirements for the host. If VMs are running apps that need high availability you'll need to consider the network requirements.

Plan host groups

Host groups act as containers for virtualization hosts and virtual machines. You apply settings at the group level, including setting resources at the host level, specifying hosts for self-service users, and storage and network options. Planning for host groups is especially important in a large-scale deployment, where host groups can help you to effectively manage resource provisioning and management.

You can base your host groups on settings that make sense for your organization. For example:

- For branch offices in your organization.
- To match your Active Directory structure.
- To reflect functions such as development, test, production, or research.
- To limit the hosts used for administrative tasks. For example you could restrict the placement of virtual machines by selecting a particular host group.
- To reserve host resources to determine the CPU, memory, disk space, disk I/O capacity, and network capacity that will always be available to the host operating system.
- To automatically place virtual machines on the most suitable host. Automatic placement is also used to deploy the virtual machines that users create in virtual machine self-service.
- To designate self-service hosts on which users can create and operate their own virtual machines. You add self-service policies to a host group to enable users or groups to create, operate, and manage their own virtual machines within a controlled environment on the hosts in the host group.

Host groups are hierarchical. For example, you can create a child host groups of a existing host group to override host reserves inherited for a parent host group, or to amend VM permissions inherited from the self-service policies of a parent host group.

- All host groups belong to the root host group - All Hosts.
- Each host or host group is identified by its host path, a sequence of host group names that specifies the location of a host or host group within the hierarchy of host groups in the navigation pane. For example, the host path All Hosts\New York\Site21\VMHost05 indicates that the host VMHost05 belongs to the host group Site21, which is a child host group of the host group New York.
- When you change the host reserves for a parent host group, you can choose whether or not to cascade the host reserve settings to hosts in all of its child host groups. If you choose to cascade the host reserve settings, all of the host reserve settings for the parent host group overwrite all previous settings for all hosts in all of the child host groups of the parent host group.
- If a parent host group is used for virtual machine self-service, each of its child host groups automatically inherits self-service policies from the parent host group. However, you can add a self-service policy for the same user or group to both a parent host group and its child host group. By adding policies to both parent and child, you can assign the same users different templates, set different virtual machine permissions, and assign a different virtual machine quota on a subset of hosts within the parent host group.
- You can use a host group to isolate a host. For example, if you have a host with guest operating systems that are running mission-critical applications, you can isolate that host by placing it in its own host group. This way you can ensure that there are no self-service policies applied on the host group and that the system resources set aside for running the host's operating system are appropriate, thus maximizing the host resources available for use by the guest operating systems.

Next steps

- [Learn about deploying the library in high availability mode](#) if required.
- [Manage the VMM library](#)

Plan the VMM networking fabric

12/7/2018 • 12 minutes to read

This article describes how to plan your networking fabric in System Center - Virtual Machine Manager (VMM).

Networking components

VMM networking contains a number of components, summarized in the following table:

NETWORKING COMPONENT	DETAILS
Logical networks	<p>In VMM your physical networks are defined as logical networks. Logical networks are a useful way of abstracting your underlying physical network infrastructure. Logical network settings will match or mirror your physical network environment. For example the IP addresses and VLAN properties will match exactly, a network site in a logical network will contain configuration settings for the site.</p> <p>By default VMM creates logical network automatically when you add a Hyper-V host to the fabric if a suitable network can't be found. You can disable this option.</p> <p>To abstract logical networks from the VMs that use them VMM provides VM networks. You connect the virtual adapter of a VM to a VM networks.</p>
MAC address pools	<p>You can create MAC address pools for VMs running on virtualization hosts in the VMM fabric. When you use static MAC address pools VMM can automatically generate and assign MAC addresses to VMs. You can use a standard pool or configure a custom pool.</p>
Load balancers	<p>VMM supports adding hardware load balancers or using NLB to load balance requests to a service tier</p>
VIP templates	<p>Virtual IP (VIP) templates contain load balancing information for a particular type of traffic. For example you could have a template that specifies how to balance HTTPS traffic on a specific load balancer.</p>
Logical switches	<p>Logical switches are containers for virtual switch settings. You apply logical switches to hosts so that you have consistent switch settings across all hosts. VMM tracks switch settings on hosts deployed with logical switches to ensure compliance.</p>

NETWORKING COMPONENT	DETAILS
Port profiles	<p>Port profiles act as containers for the properties you want a network adapter to have. Instead of configuring properties per network adapter you set up in the port profile and apply that profile to an adapter.</p> <p>There are two types of port profiles. Virtual port profiles contain settings that are applied to virtual network adapters connect to VMs or used by virtualization hosts. Uplink port profiles are used to define how a virtual switch connects to a logical network.</p>
Port classifications	<p>Port classifications are abstract containers for virtual port profile settings. This abstraction means that admins and tenants can assign a port classification to a VM template, while the VM's logical switch determines which port profile should be used. rofiles and then both admins and tenant can select a suitable classification. VMM contains a number of default port classifications. For example there's a classification for VMs that need high bandwidth and a different one for VMs that need low bandwidth. Port classifications are linked to virtual port profiles when you configure logical switches.</p>

Plan logical networks

During deployment you'll need to create logical networks and set up network sites and IP addressing in each network. Then you'll create VM networks based on those logical networks.

Here's what you'll need to plan:

1. **Automation creation:** Decide whether you want to let VMM create logical networks. VMM will automatically create a logical network each time you add a virtualization host. VMM doesn't create network sites in the automatically-created logical network. You can turn this option off in **Settings > General > Network Settings** and clear **Automatic creation of logical networks**.
2. **Logical network capacity:** If you're going to create logical networks manually figure out what you'll need to represent your physical network topology. For example if you need a management network and a network used by VMs you should create two logical networks.
3. **Logical network types:** Figure out the type of logical network you need. You'll configure VM networks on top of logical networks and those VM networks can provide network virtualization with the ability to crate multiple virtual networks on a shared physical networks, or VM networks can provide isolation with VLANs and PVLANS. When you configure the logical network you'll need to indicate the type of network you need.
4. **Network sites:** Determine how many network sites you need in the logical network. You could plan around host groups and host locations. For example a Seattle host group and a New York host group. You don't need network sites if you don't have VLANs and you're using DHCP to allocate IP addresses.
5. **VLANs/subnets:** Figure out the VLANs and IP subnets you need in the logical network. These will mirror what you have in your physical network topology.
6. **IP addressing:** If you're using static IP address assignment determine which logical networks need static address pools.

Here's what you'll need to do:

1. Identify baseline logical networks: Identify a set of initial logical networks that mirror the physical networks in your environment.
2. Identify additional logical networks for specific requirements: Define logical networks with specific purpose or perform a particular function within your environment. One of the benefits of logical networks is that you can

separate computer and network services with different business purposes without needing to change your physical infrastructure.

3. Determine isolation requirements: Identify which logical networks need to be isolated and how that isolation will be enforced, either through physical separation, VLAN/PVLAN, or network virtualization. Remember that you need isolation if the logical network is going to be used by multiple tenants. If you have a single tenant or customer isolation is optional. In turn, if you don't need isolation you'll only need a single VM network that maps to the logical network.
4. Determine the network sites, VLANs, PVLANs, and IP pools that need to be defined for each logical network you have identified.
5. Figure out which logical network will associate with which virtualization hosts.

Plan logical networks, network sites, and IP address pools

Use the following table to plan for the logical networks, VM networks, and IP address pools you will need to support a virtualized infrastructure.

ITEM TO REVIEW OR DETERMINE	DESCRIPTION AND (AS NEEDED) LINKS WITHIN THIS TOPIC
Logical networks already created by default by VMM	When you add a Hyper-V host to VMM, logical networks may be created by default, based on DNS suffixes.
How many logical networks you need, and the purpose of each	Plan to create logical networks to represent the network topology for your hosts. For example, if you need a management network, a network used for cluster heartbeats, and a network used by virtual machines, create a logical network for each.
Categories that your logical networks fall into	<p>Review the purposes of your logical networks, and categorize them:</p> <ul style="list-style-type: none">- No isolation: For example, a cluster-heartbeat network for a host cluster.- VLAN: Isolation provided by your VLANs.- Virtualized: Provides a foundation for Hyper-V network virtualization.- External: Managed through a network manager (vendor network-management console or virtual switch extension manager) outside of VM.- IPAM: Managed through an IP Address Management (IPAM) server.
How many network sites are needed in each logical network	One common way to plan network sites is around host groups and host locations. For example, for a "Seattle" host group and a "New York" host group, if you had a MANAGEMENT logical network, you might create two network sites, called MANAGEMENT - Seattle and MANAGEMENT - New York .
Which VLANs and/or IP subnets are needed in each network site	The VLANs and IP subnets you assign should match your topology.
Which logical networks (or specifically, which network sites) will need IP address pools	Determine which logical networks will use static IP addressing or load balancing, and which logical networks will be the foundation for network virtualization. For these logical networks, plan for IP address pools.

Logical networks created by default

In the VMM console, **Fabric > Networking > Logical networks**, you might see logical networks created by VMM by default. VMM creates these networks to ensure that when you add a host, you have at least one logical network for deploying virtual machines and services. No network sites are created automatically.

To illustrate how these settings work, suppose that you have not changed the settings, and you add a Hyper-V host to VMM management. In this case, VMM automatically creates logical networks that match the first DNS suffix label of the connection-specific DNS suffix on each host network adapter. On the logical network, VMM also creates a VM network that is configured with "no isolation." For example, if the DNS suffix for the host network adapter was corp.contoso.com, VMM would create a logical network named "corp," and on it, a VM network named "corp" that is configured with no isolation.

Guidelines for network sites: VLAN and IP subnet settings

The main guideline specifying VLANs and IP subnets for network sites is to reflect your network topology. For details, see the following table.

NOTE

Network sites are sometimes referred to as "logical network definitions," for example, in Windows PowerShell commands.

PURPOSE OF LOGICAL NETWORK	GUIDELINE FOR NETWORK SITES IN THAT LOGICAL NETWORK
Static IP: Logical network that will use static IP addressing, for example, a network that supports host cluster nodes	Create at least one network site and associate at least one IP subnet with the network site.
DHCP (but not VLANs): Logical network that does not include VLANs, with all computers or devices using DHCP	No network sites are necessary.
VLANs: Logical network for VLAN-based independent networks	<ul style="list-style-type: none">- If the VLANs use static IP addressing, create corresponding network sites that specify VLAN and IP subnet information.- If the VLANs use DHCP, create corresponding network sites that specify only VLAN information (no subnets).
Network virtualization: Logical network that will be the foundation for VM networks using network virtualization	<p>Create at least one network site and associate at least one IP subnet with the site. The IP subnet is necessary because this logical network will need an IP address pool.</p> <p>Assign a VLAN to the network site if appropriate.</p>
Load balancing: Logical network that will include a load balancer that is managed by VMM	Create at least one network site and associate at least one IP subnet with the network site.

NOTE

For an external network, that is, a network managed through a vendor network-management console or virtual switch extension manager outside of VMM, you can configure settings through the vendor network-management console, and allow them to be imported from the vendor network-management database into VMM.

Guidelines for IP address pools

In general, create IP address pools where you will use static IP addressing or load balancing; also create IP address pools on logical networks that will be the foundation for VM networks supporting network virtualization. VMM

uses IP address pools to assign IP addresses to Hyper-V hosts that you deploy through VMM, and to Windows-based virtual machines that you deploy through VMM, regardless of the type of host they are running on (Hyper-V or VMware ESX).

The following table provides detailed guidelines. Additional information about IP address pools is provided after the table.

PURPOSE OF LOGICAL NETWORK	GUIDELINE FOR CREATING IP ADDRESS POOLS FOR THAT LOGICAL NETWORK, OR FOR VM NETWORKS BUILT ON THAT LOGICAL NETWORK
Static IP: Logical network with "no isolation," and requiring static IP addressing, for example, a network that supports host cluster nodes	<p>Create one or more IP address pools for the logical network.</p> <p>For a logical network with "no isolation," if you create a VM network on the logical network, any IP address pools will automatically become available on the VM network. In other words, the VM network will give direct access to the logical network.</p>
VLANs: Logical network for VLAN-based independent networks, using static IP addressing (rather than DHCP)	<p>Create IP address pools on the logical network—one IP address pool for each VLAN where static IP addressing will be used.</p> <p>Later, when you create the VM networks that represent the VLANs, the IP address pools will automatically become available on those VM networks.</p>
Network virtualization: Logical network that will be the foundation for VM networks using network virtualization	<p>Create IP address pools on the logical network that provides the foundation for the VM networks. Later, when you create the VM networks, you will also create IP address pools on them (and see the important note after this table). If you use DHCP on the VM networks, VMM will respond to a DHCP request with an address from an IP address pool.</p> <p>The process of creating an IP address pool for a VM network is similar to the process of creating an IP address pool for a logical network.</p>
Load balancing: Logical network that will be the foundation for a VM network, where you will use load balancing in a "service tier" (part of a set of virtual machines deployed together as a VMM "service")	<p>Create a static IP address pool on the VM network, and in it, define a reserved range of IP addresses. When you use VMM to deploy a load-balanced service tier that uses the VM network, VMM uses the reserved range of IP addresses to assign virtual IP (VIP) addresses to the load balancer.</p>

IMPORTANT

If you configure a virtual machine to obtain a static IP address from an IP address pool, you must also configure the virtual machine to use a static MAC address. You can either specify the MAC address manually (during the **Configure Settings** step) or have VMM automatically assign a MAC address from a MAC address pool.

This requirement for static MAC addresses is necessary because VMM uses the MAC address to identify which network adapter to set the static IP address to, and this identification must happen before the virtual machine starts. Identifying the network adapter is especially important if a virtual machine has multiple adapters. If the MAC addresses were assigned dynamically through Hyper-V, VMM could not consistently identify the correct adapter to set a static IP address on.

VMM provides static MAC address pools by default, but you can customize the pools. VMM provides static MAC address pools by default, but you can customize the pools.

- When you create a static IP address pool, you can configure associated information, such as default gateways, Domain Name System (DNS) servers, DNS suffixes, and Windows Internet Name Service

(WINS) servers. All of these settings are optional.

- IP address pools support both IPv4 and IPv6 addresses. However, you cannot mix IPv4 and IPv6 addresses in the same IP address pool.

NOTE

After a virtual machine has been deployed in VMM, you can view the IP address or addresses assigned to that virtual machine. To do this, right-click the listing for the virtual machine, click **Properties**, click the **Hardware Configuration** tab, click the network adapter, and in the results pane, click **Connection details**.

Next steps

- [Set up the networking fabric](#)

Supported storage arrays

8/22/2019 • 2 minutes to read

Virtualized workloads in System Center - Virtual Machine Manager (VMM) need storage resources that meet capacity and performance requirements. VMM recognizes local and remote storage. It supports storage on block-level storage devices that expose logical unit numbers (LUNs) using fiber channel, iSCSI, and SAS connections, and network file shares.

This article provides a list of supported storage devices for VMM 2016, 1801, 1807.

DEVICE	PROTOCOL	MIN CONTROLLER FIRMWARE	SMI-S	DETAILS
Hewlett Packard Enterprise 3PAR	SMI-S	3PAR: 3PAR v. 3.2.2 MU3 or later 3PAR 8000 & 20000, 7000 & 10000	SMI-S CIM version 1.5	Link
Tintri VMstore	SMB	4.2 and later	Embedded version 2.1	Support
NEC / NEC Storage M-Series M310, M510, M710, M310F, M710F, M320, M320F	iSCSI/FC	Min Controller Firmware: 010A(Storage Control Software Revision 0941 or later)	SMI-S v1.6.1	Details
Fujitsu/ETERNUS DX60 S3, DX60 S4, DX100 S3, DX100 S4, DX200 S3, DX200 S4, DX500 S3, DX500 S4, DX600 S3, DX600 S4, DX8700 S3, DX8900 S3, DX8900 S4, DX200F, AF250, AF250 S2, AF650, AF650 S2	iSCSI/FC	V10L60 or later	EMBEDDED SMI-S v1.6.0	Storage System ETERNUS
DELL-EMC XtremIO All Flash Array	iSCSI/FC	XtremIO XMS Server Versions: 4.2.0 and 4.2.1	Embedded SMI-S V1.6.1	Link
DELL SC Series	iSCSI/FC	SCOS: 7.2 or later DSM: 2016 R3 or later SCOS 7.3.5 or later, DSM 2018 R1.20 or later	SMI-S version 1.6	Link

DEVICE	PROTOCOL	MIN CONTROLLER FIRMWARE	SMI-S	DETAILS
NetApp FAS	iSCSI/FC/SMB	8.2 and later	Proxy NetApp SMIS Provider 5.2.4 or later	Link
Huawei OceanStor V3 Series	iSCSI/FC	V300R006 and later	Huawei SMI-S version 2.1.01 or later	Link
Huawei OceanStor Dorado V3 Series	iSCSI/FC	V300R001 and later	Huawei SMI-S version 2.1.01 or later	Link
Huawei OceanStor V5 Series	iSCSI/FC	V500R007 and later	Huawei SMI-S Provider v 2.1.03 and later	TBA
IBM XIV Storage System Gen3	iSCSI/FC	11.6.0 and later	embedded SMI-S v1.6.1	Link
IBM FlashSystem A9000/A9000R	iSCSI/FC	12.1.0 and later	embedded SMI-S v1.6.1	Link
IBM Spectrum Virtualize Family Product - SAN Volume Controller(SVC)/Storwize V7000, Storwize V5000, IBM FlashSystem V9000	iSCSI/FC	7.8.0 and later	embedded SMI-S v1.6.1	Link
Pure Storage FlashArray M	iSCSI/FC	Purity 4.9 and later	SMI-S v1.6.1	Link
DELL-EMC PS Series	iSCSI	PS Firmware: v8.1 or later HIT Kit: v5.1 or later	SMP v5.1	Link

This article details the supported arrays for System Center 2019 - Virtual Machine Manager (VMM).

Supported storage arrays in VMM 2019

Virtualized workloads in System Center - Virtual Machine Manager (VMM) need storage resources that meet capacity and performance requirements. VMM recognizes local and remote storage. It supports storage on block-level storage devices that expose logical unit numbers (LUNs) using fiber channel, iSCSI, and SAS connections, and network file shares.

DEVICE	PROTOCOL	MIN CONTROLLER FIRMWARE	SMI-S	DETAILS
Hewlett Packard Enterprise 3PAR	SMI-S	3PAR: 3PAR v. 3.2.2 MU3 or later 3PAR 8000 & 20000, 7000 & 10000	SMI-S CIM version 1.5	Link
NEC / NEC Storage M-Series M510, M710, M710F, M320, M320F	iSCSI/FC	010A M510/M710 Storage Control Software revision 0973 or later M320/M320F Storage Control Software revision 1028 or later	SMI-S 1.6.1	Link

Next steps

- [Learn more](#) about configuring storage in the VMM fabric.
- Learn more about array SMI-S [Conformance Testing Program](#)

Upgrade System Center Virtual Machine Manager

3/14/2019 • 42 minutes to read

This article provides the upgrade information for System Center 2019 - Virtual Machine Manager (VMM).

Upgrade to System Center 2019 - Virtual Machine Manager

The following sections provide information about how to upgrade to VMM 2019. These include prerequisites, upgrade instructions, and tasks to complete after the upgrade finishes.

Requirements and limitations

- You should be running VMM on System Center 2016 update rollup 6 or 1801 or 1807.
- Ensure that the server meets all requirements for VMM 2019, and that prerequisites are in place [Learn more](#)
- Make sure you're running a [supported version of SQL Server](#).
- If your current VMM deployment is integrated with Azure Site Recovery, note that:
 - Site Recovery settings can't be upgraded. After the upgrade you need to redeploy.
 - Verify [Hyper-V host support](#) for VMM 2019.

Before you start

Ensure the following:

1. Complete any jobs that are currently running in VMM. note that the jobs history is deleted during the upgrade.
2. Close any connections to the VMM management server, including the VMM console and the VMM command shell.
3. Close any other programs that are running on the VMM management server.
4. Ensure that there are no pending restarts on VMM servers.
5. Perform a full backup of the VMM database.
6. If the current SQL Server database used Always On availability groups:
 - If the VMM database is included in the availability group, remove it in SQL Server Management Studio.
 - Initiate a failover to the computer that is running SQL Server, on which the VMM database is installed.
7. If you're running Operations Manager with VMM, disconnect the connection between VMM and Operations Manager server.

Upgrade sequence for System Center components

If you're running more than one System Center component, they should be upgraded in a specific order as shown below:

1. Service Management Automation
2. Orchestrator
3. Service Manager
4. Data Protection Manager

5. Operations Manager
6. Configuration Manager
7. Virtual Machine Manager
8. Service Provider Foundation

Upgrade a standalone VMM server

NOTE

When you are upgrading a standalone VMM server, we recommend that you install VMM 2019 on the same server that had VMM 2016, 1801 or 1807.

If you are using Distributed Key Management, you may choose to install VMM 2019 on a different server but make sure that the new server has the same name as that of the old VMM server.

Use the following procedures:

- [Back up and upgrade the OS](#)
- [Install VMM 2019](#)

Back up and upgrade OS

1. Back up and retain the VMM database.
2. [Uninstall the VMM](#). Ensure to remove both the management server and console.
3. Upgrade the management OS to Windows Server 2019.
4. Install Windows 10 version of ADK.

Uninstall the VMM

1. Go to **Control Panel > Programs > Program and Features**, select **Virtual Machine Manager** and click **Uninstall**.
2. On the **Uninstall wizard**, select **Remove Features**, select both **VMM management Server** and **VMM Console** under the features to remove, list.
3. On database options page, select **Retain database**.
4. Review the summary and click **Uninstall**.

Install VMM 2019

1. In the main setup page, click **Install**.
2. In **Select features to install**, select the VMM management server, and then click **Next**. The VMM console will be automatically installed.
3. In **Product registration information**, provide the appropriate information, and then click **Next**. If you do not enter a product key, VMM will be installed as an evaluation version that expires after 180 days from the installation date.
4. In **Please read this license agreement**, review the license agreement, select the **I have read, understood, and agree with the terms of the license agreement** check box, and then click **Next**.
5. In **Usage and Connectivity Data**, select either of the options, and then click **Next**.
6. If **Microsoft Update** page appears, select whether you want to use Microsoft Update and click **Next**. If you've already chosen to use Microsoft Update on this computer, this page won't appear.
7. In **Installation location**, use the default path or type a different installation path for the VMM program files, and then click **Next**.

8. In **Database configuration**:

- [Learn more](#) if you need to upgrade the VMM SQL Server.
- If you're using a remote SQL instance, specify the SQL server computer name.
- If SQL server runs on the VMM server, type the name of the VMM server, or type **localhost**. If the SQL Server is in a cluster, type the cluster name.
- Don't specify a port value if you're using local SQL server, or if your remote SQL server uses the default port (1443).
- Select **Existing Database** and select the database that you retained (backed up) from your previous installation. Provide credentials with permissions to access the database. When you're prompted to upgrade the database click **Yes**.

9. In **Configure service account and distributed key management**, specify the account that the VMM service will use.

NOTE

You can't change the identity of the VMM service account after installation.

10. Under **Distributed Key Management**, select whether to store encryption keys in Active Directory.

NOTE

Choose the settings for the service account and distributed key management, carefully. Based on your selection, encrypted data such as passwords in templates might not be available after the upgrade, and you'll need to enter them manually.

11. In **Port configuration**, use the default port number for each feature or provide a unique port number that is appropriate in your environment.

NOTE

You cannot change the ports that you assign during the installation of a VMM management server unless you uninstall and then reinstall the VMM management server. Also, do not configure any feature to use port 5986, this port number is preassigned.

12. In **Library configuration**, select whether to create a new library share or to use an existing library share on the computer. The default library share that VMM creates is named **MSSCVMMLibrary**, and the folder is located at **%SYSTEMDRIVE%\ProgramData\Virtual Machine Manager Library Files**. **ProgramData** is a hidden folder, and you cannot remove it. After the VMM management server is installed, you can add library shares and library servers by using the VMM console or by using the VMM command shell.

13. In **Upgrade compatibility report**, review the settings, click **Next** to proceed with the upgrade.

14. In **Installation Summary**, review the settings and click **Install** to upgrade the server. **Installing features** page appears and displays the installation progress.

15. In **Setup completed successfully**, click **Close** to finish the installation. To open the VMM console, check **Open the VMM console when this wizard closes**, or you can click the Virtual Machine Manager Console icon on the desktop.

NOTE

Once the upgrade is successful, [upgrade the host agent manually](#), by using the VMM.

In case of any issue with the setup, check the **%SYSTEMDRIVE%\ProgramData\VMMLogs** folder.

During the setup, VMM enables the following firewall rules. These rules remain in effect even if you uninstall the VMM later.

- Windows Remote Management
- Windows Standards-Based Storage Management

Upgrade a highly available VMM server

You can upgrade a highly available (HA) VMM server 2016 or 1801 or 1807 to 2019.

The following two modes of upgrade are supported:

- [Mixed mode with no additional VMM servers](#)
- [Mixed mode with additional VMM servers](#)

NOTE

SQL Server upgrade can be performed any time, independent of VMM upgrade.

Mixed mode upgrade with no additional VMM servers

This procedure requires no additional VMM servers, but has increased risk for downtime in some scenarios. For example, when you have two node HA VMM and the active VMM node fails while you are upgrading the passive. In this scenario, your VMM server will not have a failover node available.

Follow these steps:

1. Backup and retain the VMM database.
2. [Uninstall the VMM](#) on the passive node.
3. On the passive VMM node, upgrade the management OS to Windows server 2019/2016.
4. Upgrade to the Windows 10 version of the ADK.
5. Install VMM 2019 on the passive node by using the following steps:
 - In the main setup page, click **Install**.
 - In **Select features to install**, select **VMM management server** and then click **Next**. The VMM console will be automatically installed.
 - When prompted, confirm that you want to add this server as a node to the highly available deployment.
 - On **Database Configuration** page, if prompted select to upgrade the database.
 - Review the summary and complete the installation.
6. Failover the active VMM node to the newly upgraded VMM server.
7. Repeat the procedure on other VMM nodes.
8. Update the cluster functional level by using the **Update-ClusterFunctionalLevel** [command](#).
9. [Optional] Install the appropriate SQL Command line utilities.

Mixed mode upgrade with additional VMM servers

This procedure requires additional VMM servers, however, ensures almost no downtime in all the scenarios.

Follow these steps:

1. Backup and retain the VMM database.
2. Add the same number of additional servers (with Windows Server 2019/2016 Management OS) that equals to the server number present in the HA cluster.
3. Install Windows 10 version of the ADK on the newly added 2019 servers.
4. Install VMM 2019 on one of the newly added servers by using the details in **step 5** in [Mixed mode upgrade with no additional VMM servers](#).
5. Repeat the installation steps for all the other newly added servers.
6. Failover the active VMM node to one of the newly added servers.
7. Remove 2016 nodes from the cluster after failover.
8. Update the cluster functional level by using the **Update-ClusterFunctionalLevel** [command](#).
9. [Optional] Install the appropriate SQL Command line utilities.

NOTE

Once the HA VMM upgrade is successful, [upgrade the host agent manually](#), by using the VMM.

Update VMM agents

After the upgrade, you need to update the VMM agents on your Hyper-V hosts and in your VMM library servers.

1. Click **Fabric > Servers > All Hosts**.
2. In the **Hosts** pane, right-click a column heading, and then click **Agent Version Status**.
3. Select the host with the VMM agent that you want to update. On the **Hosts** tab, in the **Host** group, click **Refresh**. If a host needs to have its VMM agent updated, the **Host Status** column will display **Needs Attention**, and the **Agent Version Status** column will display **Upgrade Available**.
4. Right-click the host with the VMM agent that you want to update, and then click **Update Agent**. In **Update Agent** provide the necessary credentials, and then click **OK**.
5. The **Agent Version Status** column will display a value of **Upgrading**. After the VMM agent is updated successfully on the host, the **Agent Version Status** column will display a value of **Up-to-date**, and the **Agent Version** column will display the updated version of the agent. After you refresh the host again, the **Host Status** column for the host will display a value of **OK**.
6. You can update the VMM agent on a VMM library server in a similar manner. To view a list of VMM library servers, click **Fabric > Servers > Library Servers**.

Reassociate hosts and library servers

After the upgrade, you might need to reassociate virtual machine hosts and VMM library servers with the VMM management server.

Follow these steps:

1. Click **Fabric > Servers > All Hosts**.
2. In the **Hosts** pane, ensure that the **Agent Status** column is displayed. If it isn't displayed, right-click a column heading, select **Agent Status**.
3. In the host group, click **Refresh**. If a host needs to be reassociated, the **Host Status** column displays **Needs Attention**, and the **Agent Status** column displays **Access Denied**. Right-click the host that you want to reassociate, and then click **Reassociate**.
4. In **Reassociate Agent** page, provide credentials, and then click **OK**. The Agent Status displays the status as **Reassociating**. After the host is reassociated successfully, the status changes to **Responding**.
6. Refresh the host, the host status columns now displays **OK**. After you have reassociated the host, you might need to update the VMM agent on the host.

Upgrade the VMM SQL Server database

There are a couple of reasons you might want to upgrade the VMM SQL Server database:

- You're upgrading VMM to System Center 2019, and the current SQL Server database version isn't supported.
- You want to upgrade a VMM standalone server to a high availability server, and SQL Server is installed locally.
- You want to move the SQL Server database to a different computer.

Collect database information

Before you upgrade, collect information about the VMM database:

1. Record the database connection in the VMM console > **Settings** > **General** > **Database Connection**.
2. Record the account information in Server Manager > **Tools** > **Services**. Right-click **System Center Virtual Machine Manager** > **Properties** > **Log On**. This is the domain or local account that was assigned as the service account when VMM was installed. You can check if it's local in **Tools** > **Computer Manager** > **Local Users and Groups** > **Users**.
3. Check whether you used distributed key management when you installed VMM, or if encryption keys are stored locally on the VMM server.
4. If you're moving the VMM database, but not upgrading VMM, check which update rollups have been applied on the VMM server.

Upgrade a standalone database

1. Back up the existing VMM database, and copy the backup to a computer running a supported version of SQL Server.
2. Use SQL Server tools to restore the database.
 - If you are upgrading VMM, you'll specify the new SQL Server location in VMM setup > **Database Configuration**.
 - If you want to upgrade the database without upgrading VMM, you need to uninstall and then reinstall VMM. When you uninstall, **Database Options** page, select **Retain database**. Then, reinstall with the same settings you used for the original installation. On the **Database Configuration** you specify the new SQL Server details. After reinstall apply the update rollups and check that the deployment is working as expected.

Upgrade a highly available database

1. Record the source version of the existing database, and the version you want to upgrade to.
2. Create a backup of the highly available SQL Server database, from the active node of the SQL Server cluster.
3. Upgrade passive SQL Server nodes to the new version. After the upgrade, optionally install SQL Server Management Studio if you want to manage SQL Server from this node.
4. Fail over the highly available SQL server role, from the currently active node to the upgraded node. After failover, you can use SQL Server Management Studio to validate the running database version.
5. Repeat the upgrade for the other nodes in the HA SQL cluster. As an additional validation, you can fail over the SQL Server database roles, to ensure that everything works as expected.

Migrate a SQL Server cluster as part of the VMM upgrade

1. Take a backup of the highly available VMM database from the active node of the existing SQL cluster.
2. Note the VMM role name to use when reinstalling the VMM server role. Uninstall VMM server from existing VMM cluster nodes with the retain database option. When uninstalling VMM server from last node, you may get a message about unsuccessful SPN registration. This is a known issue with no functional impact.
3. Restore the backed-up DB into another SQL cluster running supported SQL version. Add the user on which VMM service is running as User to this new DB with membership to db_owner.
4. While upgrading VMM Server as part of SQL Cluster migration, give the Parameters corresponding to new SQL Cluster.

Redeploy Azure Site Recovery

If Azure Site Recovery was integrated into your VMM 2016/1801/1807 deployment, you need to redeploy it with VMM 2019, for [replication to Azure](#), or [replication to a secondary site](#).

Connect to Operations Manager

After the upgrade, reconnect VMM to Operations Manager.

Renew certificates for PXE servers

If you have a PXE server in the VMM fabric, you need to remove it from the fabric, and then add it again. This is in order to renew the PXE server certificate, and avoid certificate errors.

You must have System Center - Virtual Machine Manager 1801 installed to apply the 1807 update. [Learn more on how to install VMM 1807](#).

This article provides the upgrade information for System Center 1801 - Virtual Machine Manager (VMM).

This article provides the upgrade information for System Center 2016 - Virtual Machine Manager (VMM).

Upgrade to System Center 1801 - Virtual Machine Manager

The following sections provide information about how to upgrade to VMM 1801. Also includes upgrade steps, and tasks you should complete after the upgrade finishes.

Requirements and limitations

- You should be running VMM on System Center 2012 R2 update rollup 14 or 2016 UR 4.
- Ensure that the server meets all requirements for VMM 1801, and that prerequisites are in place [Learn more](#)
- Make sure you're running a [supported version of SQL Server](#).
- If your current VMM deployment is integrated with Azure Site Recovery, note that:
 - Site Recovery settings can't be upgraded. After the upgrade you need to redeploy.
 - Verify [Hyper-V host support](#) for VMM 1801.

Before you start

Ensure the following:

1. Complete any jobs that are currently running in VMM. note that the jobs history is deleted during the upgrade.
2. Close any connections to the VMM management server, including the VMM console and the VMM command shell.
3. Close any other programs that are running on the VMM management server.
4. Ensure that there are no pending restarts on VMM servers.
5. Perform a full backup of the VMM database.
6. If you're running Operations Manager with VMM, disconnect the connection between VMM and Operations Manager server.

Upgrade sequence for System Center components

If you're running more than one System Center component, they should be upgraded in a specific order:

1. Service Management Automation
2. Orchestrator

3. Service Manager
4. Data Protection Manager (DPM)
5. Operations Manager
6. Configuration Manager
7. Virtual Machine Manager (VMM)
8. App Controller
9. Service Provider Foundation
10. Windows Azure Pack for Windows Server
11. Service Bus Clouds
12. Windows Azure Pack
13. Service Reporting

Upgrade a standalone VMM server

Use the following procedures:

- [Back up and upgrade the OS](#)
- [Install VMM 1801](#)

Back up and upgrade OS

1. Back up and retain the VMM database.
2. [Uninstall the VMM](#). Ensure to remove both the management server and console.
3. Upgrade the management OS to Windows Server 2016/1709.
4. Install Windows 10 or 1709 version of ADK respectively.

Uninstall the VMM

1. Go to **Add remove programs**, select **VMM** and click **Uninstall**.
2. On the **Uninstall wizard**, select **Remove Features**, select both **VMM management Server** and **VMM Console** under the features to remove, list.
3. On database options page, select **Retain database**.
4. Review the summary and click **Uninstall**.

Install VMM 1801

1. In the main setup page, click **Install**.
2. In **Select features to install**, select the VMM management server and then click **Next**. The VMM console will be automatically installed.
3. In **Product registration information**, provide the appropriate information, and then click **Next**. If you do not enter a product key, VMM will be installed as an evaluation version that expires after 180 days from the installation date.
4. In **Please read this license agreement**, review the license agreement, select the **I have read, understood, and agree with the terms of the license agreement** ****check box**, and then click ****Next**.
5. In **Usage and Connectivity Data**, select either of the options, and then click **Next**.
6. If **Microsoft Update** page appears, select whether you want to use Microsoft Update and click **Next**. If you've already chosen to use Microsoft Update on this computer, this page won't appear.
7. In **Installation location**, use the default path or type a different installation path for the VMM program files, and then click **Next**.
8. In **Database configuration**:

- [Learn more](#) if you need to upgrade the VMM SQL Server.
- If you're using a remote SQL instance, specify the SQL server computer name.
- If SQL server runs on the VMM server, type the name of the VMM server, or type **localhost**. If the SQL Server is in a cluster, type the cluster name.
- Don't specify a port value if you're using local SQL server, or if your remote SQL server uses the default port (1443).
- Select **Existing Database** and select the database that you retained (backed up) from your previous installation. Provide credentials with permissions to access the database. When you're prompted to upgrade the database click **Yes**.

9. In **Configure service account and distributed key management**, specify the account that the VMM service will use.

NOTE

You can't change the identity of the VMM service account after installation.

10. Under **Distributed Key Management**, select whether to store encryption keys in Active Directory.

NOTE

Choose the settings for the service account and distributed key management, carefully. Based on your selection, encrypted data such as passwords in templates might not be available after the upgrade, and you'll need to enter them manually.

11. In **Port configuration**, use the default port number for each feature or provide a unique port number that is appropriate in your environment.

NOTE

You cannot change the ports that you assign during the installation of a VMM management server unless you uninstall and then reinstall the VMM management server. Also, do not configure any feature to use port 5986, this port number is preassigned.

12. In **Library configuration**, select whether to create a new library share or to use an existing library share on the computer. The default library share that VMM creates is named **MSSCVMMLibrary**, and the folder is located at **%SYSTEMDRIVE%\ProgramData\Virtual Machine Manager Library Files**. **ProgramData** is a hidden folder, and you cannot remove it. After the VMM management server is installed, you can add library shares and library servers by using the VMM console or by using the VMM command shell.
13. In **Upgrade compatibility report**, review the settings, click **Next** to proceed with the upgrade.
14. In **Installation Summary**, review the settings and click **Install** to upgrade the server. **Installing features** page appears and displays the installation progress.
15. In **Setup completed successfully**, click **Close** to finish the installation. To open the VMM console, check **Open the VMM console when this wizard closes**, or you can click the Virtual Machine Manager Console icon on the desktop.

NOTE

Once the upgrade is successful, [upgrade the host agent manually](#), by using the VMM.

In case of any issue with the setup, check the **%SYSTEMDRIVE%\ProgramData\VMMLogs** folder.

During the setup, VMM enables the following firewall rules. These rules remain in effect even if you uninstall the VMM later.

- Windows Remote Management
- Windows Standards-Based Storage Management

Upgrade a highly available VMM server

You can upgrade a highly available (HA) VMM server 2012 R2 and 2016 to 1801.

The following two modes of upgrade are supported:

- [Mixed mode with no additional VMM servers](#)
- [Mixed mode with additional VMM servers](#)

NOTE

SQL Server upgrade can be performed any time, independent of VMM upgrade.

Mixed mode upgrade with no additional VMM servers

This procedure requires no additional VMM servers, but has increased risk for downtime in some scenarios. For example, when you have two node HA VMM and the active VMM node fails while you are upgrading the passive. In this scenario, your VMM server will not have a failover node available.

Follow these steps:

1. Backup and retain the VMM database.
2. [Uninstall the VMM](#) on the passive node.
3. VMM 1801 supports WS 2016 and 1709 as Management OS. On the passive VMM node, upgrade the management OS to Windows server 2016/1709.

If the cluster has VMM 2012 R2 on Windows Server (WS) 2012 R2 and you want to upgrade to VMM 1801 on WS 1709, then use the following steps for each of such nodes in the cluster:

- Upgrade to VMM 1801 on WS 2016 as the management OS for all the cluster nodes with WS 2012 R2. This is because, 2012 R2 and WS 1709 mixed cluster is not supported.
 - Upgrade the management OS to WS 1709.
4. Upgrade to the Windows 10/1709 version of the ADK.
 5. Install VMM 1801 on the passive node by using the following steps:
 - In the main setup page, click **Install**.
 - In **Select features to install**, select **VMM management server** and then click **Next**. The VMM console will be automatically installed.
 - When prompted, confirm that you want to add this server as a node to the highly available deployment.
 - On **Database Configuration** page, if prompted select to upgrade the database.
 - Review the summary and complete the installation.
 6. Failover the active VMM node to the newly upgraded VMM server.
 7. Repeat the procedure on other VMM nodes.
 8. Update the cluster functional level by using the **Update-ClusterFunctionalLevel** [command](#).

9. [Optional] Install the appropriate SQL Command line utilities.

Mixed mode upgrade with additional VMM servers

This procedure requires additional VMM servers, however, ensures almost no downtime in all the scenarios.

Follow these steps:

1. Backup and retain the VMM database.
2. Add the same number of additional servers (with Windows Server 2016/1709 Management OS) that equals to the server number present in the HA cluster.

If the cluster has VMM 2012 R2 nodes on WS 2012 R2 and you want to upgrade to WS 1709 with VMM 1801, then use the following steps:

- Add new WS 2016 nodes and install VMM 1801 on all of these new nodes.
 - Remove all the WS 2012 R2 nodes from the cluster.
 - Upgrade the management OS to WS 1709.
3. Install Windows 10/1709 version of the ADK on the newly added 2016 servers.
 4. Install VMM 1801 on one of the newly added servers by using the details in **step 5** in [Mixed mode upgrade with no additional VMM servers](#mixed-mode- upgrade-with-no-additional-VMM-servers).
 5. Repeat the installation steps for all the other newly added servers.
 6. Failover the active VMM node to one of the newly added servers.
 7. Remove 2012 R2/2016 nodes from the cluster after failover.
 8. Update the cluster functional level by using the **Update-ClusterFunctionalLevel** [command](#).
 9. [Optional] Install the appropriate SQL Command line utilities.

NOTE

Once the HA VMM upgrade is successful, [upgrade the host agent manually](#), by using the VMM.

Upgrade the VMM SQL Server database

There are a couple of reasons you might want to upgrade the VMM SQL Server database:

- You're upgrading VMM to System Center 1801, and the current SQL Server database version isn't supported.
- You want to upgrade a VMM standalone server to a high availability server, and SQL Server is installed locally.
- You want to move the SQL Server database to a different computer.

Collect database information

Before you upgrade, collect information about the VMM database:

1. Record the database connection in the VMM console > **Settings** > **General** > **Database Connection**.
2. Record the account information in Server Manager > **Tools** > **Services**. Right-click **System Center Virtual Machine Manager** > **Properties** > **Log On**. This is the domain or local account that was assigned as the service account when VMM was installed. You can check if it's local in **Tools** > **Computer Manager** > **Local Users and Groups** > **Users**.
3. Check whether you used distributed key management when you installed VMM, or if encryption keys are stored locally on the VMM server.
4. If you're moving the VMM database, but not upgrading VMM, check which update rollups have been applied on the VMM server.

Upgrade a standalone database

1. Back up the existing VMM database, and copy the backup to a computer running a supported version of SQL Server.
2. Use SQL Server tools to restore the database.
 - If you are upgrading VMM, you'll specify the new SQL Server location in VMM setup > **Database Configuration**.
 - If you want to upgrade the database without upgrading VMM, you need to uninstall and then reinstall VMM. When you uninstall, **Database Options** page, select **Retain database**. Then, reinstall with the same settings you used for the original installation. On the **Database Configuration** you specify the new SQL Server details. After reinstall apply the update rollups and check that the deployment is working as expected.

Upgrade a highly available database

1. Record the source version of the existing database, and the version you want to upgrade to.
2. Create a backup of the highly available SQL Server database, from the active node of the SQL Server cluster.
3. Upgrade passive SQL Server nodes to the new version. After the upgrade, optionally install SQL Server Management Studio if you want to manage SQL Server from this node.
4. Fail over the highly available SQL server role, from the currently active node to the upgraded node. After failover, you can use SQL Server Management Studio to validate the running database version.
5. Repeat the upgrade for the other nodes in the HA SQL cluster. As an additional validation, you can fail over the SQL Server database roles, to ensure that everything works as expected.

Migrate a SQL Server cluster as part of the VMM upgrade

1. Take a backup of the highly available VMM database from the active node of the existing SQL cluster.
2. Note the VMM role name to use when reinstalling the VMM server role. Uninstall VMM server from existing VMM cluster nodes with the retain database option. When uninstalling VMM server from last node, you may get a message about unsuccessful SPN registration. This is a known issue with no functional impact.
3. Restore the backed-up DB into another SQL cluster running supported SQL version. Add the user on which VMM service is running as User to this new DB with membership to db_owner.
4. While upgrading VMM Server as part of SQL Cluster migration, give the Parameters corresponding to new SQL Cluster.

Update VMM agents

After the upgrade, you need to update the VMM agents on your Hyper-V hosts and in your VMM library servers.

1. Click **Fabric > Servers > All Hosts**.
2. In the **Hosts** pane, right-click a column heading, and then click **Agent Version Status**.
3. Select the host with the VMM agent that you want to update. On the **Hosts** tab, in the **Host** group, click **Refresh**. If a host needs to have its VMM agent updated, the **Host Status** column will display **Needs Attention**, and the **Agent Version Status** column will display **Upgrade Available**.
4. Right-click the host with the VMM agent that you want to update, and then click **Update Agent**. In **Update Agent** provide the necessary credentials, and then click **OK**.
5. The **Agent Version Status** column will display a value of **Upgrading**. After the VMM agent is updated successfully on the host, the **Agent Version Status** column will display a value of **Up-to-date**, and the **Agent Version** column will display the updated version of the agent. After you refresh the host again, the **Host Status** column for the host will display a value of **OK**.
6. You can update the VMM agent on a VMM library server in a similar manner. To view a list of VMM library servers, click **Fabric > Servers > Library Servers**.

Redeploy Azure Site Recovery

If Azure Site Recovery was integrated into your VMM 2012 R2/2016 deployment, you need to redeploy it with VMM 1801, for [replication to Azure](#), or [replication to a secondary site](#).

Read this [blog entry](#) for details of Hyper-V host support when running VMM 1801.

Connect to Operations Manager

After the upgrade, reconnect VMM to Operations Manager.

Renew certificates for PXE servers

If you have a PXE server in the VMM fabric, you need to remove it from the fabric, and then add it again. This is in order to renew the PXE server certificate, and avoid certificate errors.

Upgrade to System Center 2016 - VMM

The following sections describe how to upgrade to System Center Virtual Machine Manager (VMM) 2016. They include prerequisites, upgrade instructions, and tasks to complete after the upgrade finishes.

Requirements and limitations

- VMM should be running on System Center 2012 R2 with update rollup 9 or later.
- The server on which you'll install VMM should meet VMM 2016 requirements. [Learn more](#)
- Make sure you're running a [supported version of SQL Server](#).
- If your current VMM deployment is integrated with Azure Site Recovery, note that Site Recovery settings can't be upgraded. After the upgrade you need to redeploy.
- Verify Hyper-V host [support requirements](#) for VMM 2016.

Before you start

1. Complete any jobs that are currently running in VMM. All job history is deleted during the upgrade.
2. Close any connections to the VMM management server, including the VMM console and the VMM command shell.
3. Close any other programs that are running on the VMM management server.
4. Ensure that there are no pending restarts on VMM servers.
5. Perform a full backup of the VMM database.
6. If the current SQL Server database used Always On availability groups:
 - If the VMM database is included in the availability group, remove it in SQL Server Management Studio.
 - Initiate a failover to the computer that is running SQL Server, and on which the VMM database is installed.
7. If you're running Operations Manager with VMM, disconnect the connection between VMM and Operations Manager server.
8. If the VMM 2012 R2 server is running update rollup 10 or 11, and you have a Citrix NetScaler load balancer deployed, run this SQL Server script before you start the upgrade, otherwise it might fail. The script isn't needed if you're running update rollup 12 or later.

```
ALTER TABLE [dbo].[tbl_NetMan_HardwareModelSettings] ALTER COLUMN Version NVARCHAR(255) NULL; GO
```

Upgrade sequence for System Center components

If you're running more than one System Center component, they should be upgraded in a specific order:

1. Service Management Automation
2. Orchestrator
3. Service Manager
4. Data Protection Manager (DPM)
5. Operations Manager
6. Configuration Manager
7. Virtual Machine Manager (VMM)
8. App Controller
9. Service Provider Foundation
10. Windows Azure Pack for Windows Server
11. Service Bus Clouds
12. Windows Azure Pack
13. Service Reporting

Upgrade a standalone VMM server

Back up and upgrade the operating system, and then install VMM 2016.

Back up and upgrade the operating system

1. Back up and retain the VMM database.
2. Uninstall VMM. To do this: a. In **Add remove programs**, select **VMM > Uninstall**. b. Select **Remove Features**, and then select **VMM management Server** and **VMM Console**. c. In **Database Options**, select **Retain database**. d. Review the summary and click **Uninstall**.
3. Upgrade the management operating system to Windows Server 2016.
4. Upgrade to the Windows 10 version of the ADK.

Install VMM 2016

1. In the main setup page, click **Install**.
2. In **Select features to install**, select the VMM management server > **Next**. The VMM console will be automatically installed.
3. In **Product registration information**, provide the appropriate information, and then click **Next**. If you do not enter a product key, VMM will be installed as an evaluation version that expires after 180 days from the installation date.
4. In **Please read this license agreement**, review the license agreement, select **I have read, understood, and agree with the terms of the license agreement**, and then click **Next**.
5. In **Usage and Connectivity Data**, select either of the options, and then click **Next**.
6. If **Microsoft Update** page appears, select whether you want to use Microsoft Update and click **Next**. If you've already chosen to use Microsoft Update on this computer, this page won't appear.
7. In **Installation location**, use the default path or type a different installation path for the VMM program files, and then click **Next**.
8. In **Database configuration**:
 - [Learn more](#) if you need to upgrade the VMM SQL Server.
 - If you're using a remote SQL instance, specify the SQL server computer name.
 - If SQL server runs on the VMM server, type the name of the VMM server, or type **localhost**. If the SQL Server is in a cluster, type the cluster name.

- Don't specify a port value if you're using local SQL server, or if your remote SQL server uses the default port (1443).
 - Select **Existing Database** and select the database that you retained (backed up) from your previous installation. Provide credentials with permissions to access the database. When you're prompted to upgrade the database click **Yes**.
9. In **Configure service account and distributed key management**, specify the account that the VMM service will use. You can't change the identity of the VMM service account after installation.
10. Under **Distributed Key Management**, select whether to store encryption keys in Active Directory.

NOTE

Choose the settings for the service account and distributed key management carefully. Based on your selection, encrypted data such as passwords in templates might not be available after the upgrade, and you'll need to enter them manually.

11. In **Port configuration**, use the default port number for each feature, or provide a unique port number that is appropriate in your environment. To change the ports you assign during installation of a VMM management server, you need to uninstall and then reinstall the server. Don't configure any feature to use port 5986, this port number is preassigned.
12. In **Library configuration**, select whether to create a new library share or to use an existing library share on the computer. The default library share that VMM creates is named **MSSCVMMLibrary**, and the folder is located at **%SYSTEMDRIVE%\ProgramData\Virtual Machine Manager Library Files**. **ProgramData** is a hidden folder, and you cannot remove it. After the VMM management server is installed, you can add library shares and library servers by using the VMM console or by using the VMM command shell.
13. In **Upgrade compatibility report**, review the settings, click **Next** to proceed with the upgrade.
14. In **Installation Summary**, review the settings and click **Install** to upgrade the server. **Installing features** page appears and displays the installation progress.
15. In **Setup completed successfully**, click **Close** to finish the installation. To open the VMM console, check **Open the VMM console when this wizard closes**, or you can click the VMM console icon on the desktop.
16. After the upgrade, [upgrade the host agent manually](#).
17. During setup, VMM enables the following firewall rules. These rules remain in effect even if you uninstall the VMM later:
- Windows Remote Management
 - Windows Standards-Based Storage Management

if you experience any issues during setup, check the logs in the **%SYSTEMDRIVE%\ProgramData\VMMLogs** folder.

Upgrade a highly available VMM server

You can upgrade a highly available VMM 2012 R2 server (running at least update rollup 9) to VMM 2016.

The following modes of upgrade are supported:

- [Mixed mode with no additional VMM servers](#)
- [Mixed mode with additional VMM servers](#)

NOTE

SQL Server upgrade can be performed any time, independent of VMM upgrade.

Mixed mode upgrade with no additional VMM servers

You don't need additional VMM servers, but has increased risk for downtime in some scenarios. For example, when you have two node HA VMM and the active VMM node fails while you are upgrading the passive. In this scenario, your VMM server will not have a failover node available.

1. Backup and retain the VMM database.
2. [Uninstall the VMM](#) on the passive node.
3. On the passive VMM node, upgrade the management OS to Windows server 2016.
4. Upgrade to the Windows 10 version of the ADK.
5. Install VMM 2016 on the passive node by using the following steps:
 - In the main setup page, click **Install**.
 - In **Select features to install**, select **VMM management server** and then click **Next**. The VMM console will be automatically installed .
 - When prompted, confirm that you want to add this server as a node to the highly available deployment.
 - On **Database Configuration** page, if prompted select to upgrade the database.
 - Review the summary and complete the installation.
6. Failover the active VMM node to the newly upgraded VMM server.
7. Repeat the procedure on other VMM nodes.
8. Update the cluster functional level by using the **Update-ClusterFunctionalLevel** [command](#).
9. [Optional] Install the appropriate SQL Command line utilities.

Mixed mode upgrade with additional VMM servers

You need additional servers. However, there's almost no downtime in all the scenarios.

1. Backup and retain the VMM database.
2. Add the same number of additional servers (with Windows Server 2016 Management OS) that equals to the server number present in the HA cluster.
3. Install Windows 10 version of the ADK on the newly added 2016 servers.
4. Install VMM 2016 on one of the newly added servers by using the details in **step 5** in [Mixed mode upgrade with no additional VMM servers](#mixed-mode- upgrade-with-no-additional-VMM-servers).
5. Repeat the installation steps for all the other newly added servers.
6. Failover the active VMM node to one of the newly added servers.
7. Remove 2012 R2 nodes from the cluster after failover.
8. Update the cluster functional level by using the [Update-ClusterFunctionalLevelcommand](#).
9. Optionally install the appropriate SQL Command line utilities.
10. After the upgrade, [upgrade the host agent manually](#).

Upgrade the VMM SQL Server database

There are a couple of reasons you might want to upgrade the VMM SQL Server database:

- You're upgrading VMM to System Center 2016, and the current SQL Server database version isn't supported.
- You want to upgrade a VMM standalone server to a high availability server, and SQL Server is installed locally.
- You want to move the SQL Server database to a different computer.

Collect database information

Before you upgrade, collect information about the VMM database:

1. Record the database connection in the VMM console > **Settings** > **General** > **Database Connection**.
2. Record the account information in Server Manager > **Tools** > **Services**. Right-click **System Center Virtual Machine Manager** > **Properties** > **Log On**. This is the domain or local account that was assigned as the service account when VMM was installed. You can check if it's local in **Tools** > **Computer Manager** > **Local Users and Groups** > **Users**.
3. Check whether you used distributed key management when you installed VMM, or if encryption keys are stored locally on the VMM server.
4. If you're moving the VMM database, but not upgrading VMM, check which update rollups have been applied on the VMM server.

Upgrade a standalone database

1. Back up the existing VMM database, and copy the backup to a computer running a supported version of SQL Server.
2. Use SQL Server tools to restore the database.
 - If you are upgrading VMM, you'll specify the new SQL Server location in VMM setup > **Database Configuration**.
 - If you want to upgrade the database without upgrading VMM, you need to uninstall and then reinstall VMM. When you uninstall, **Database Options** page, select **Retain database**. Then, reinstall with the same settings you used for the original installation. On the **Database Configuration** you specify the new SQL Server details. After reinstall apply the update rollups and check that the deployment is working as expected.

Upgrade a highly available database

1. Record the source version of the existing database, and the version you want to upgrade to.
2. Create a backup of the highly available SQL Server database, from the active node of the SQL Server cluster.
3. Upgrade passive SQL Server nodes to the new version. After the upgrade, optionally install SQL Server Management Studio if you want to manage SQL Server from this node.
4. Fail over the highly available SQL server role, from the currently active node to the upgraded node. After failover, you can use SQL Server Management Studio to validate the running database version.
5. Repeat the upgrade for the other nodes in the HA SQL cluster. As an additional validation, you can fail over the SQL Server database roles, to ensure that everything works as expected.

Migrate a SQL Server cluster as part of the VMM upgrade

1. Take a backup of the highly available VMM database from the active node of the existing SQL cluster.
2. Note the VMM role name to use when reinstalling the VMM server role. Uninstall VMM server from existing VMM cluster nodes with the retain database option. When uninstalling VMM server from last node, you may get a message about unsuccessful SPN registration. This is a known issue with no functional impact.
3. Restore the backed-up DB into another SQL cluster running supported SQL version. Add the user on which VMM service is running as User to this new DB with membership to db_owner.
4. While upgrading VMM Server as part of SQL Cluster migration, give the Parameters corresponding to new SQL Cluster.

Update VMM agents

After the upgrade, you need to update the VMM agents on your Hyper-V hosts and in your VMM library servers.

1. Click **Fabric** > **Servers** > **All Hosts**.
2. In the **Hosts** pane, right-click a column heading, and then click **Agent Version Status**.
3. Select the host with the VMM agent that you want to update. On the **Hosts** tab, in the **Host** group, click **Refresh**. If a host needs to have its VMM agent updated, the **Host Status** column will display **Needs**

Attention, and the **Agent Version Status** column will display **Upgrade Available**.

4. Right-click the host with the VMM agent that you want to update, and then click **Update Agent**. In **Update Agent** provide the necessary credentials, and then click **OK**.
5. The **Agent Version Status** column will display a value of **Upgrading**. After the VMM agent is updated successfully on the host, the **Agent Version Status** column will display a value of **Up-to-date**, and the **Agent Version** column will display the updated version of the agent. After you refresh the host again, the **Host Status** column for the host will display a value of **OK**.
6. You can update the VMM agent on a VMM library server in a similar manner. To view a list of VMM library servers, click **Fabric > Servers > Library Servers**.

Reassociate hosts and library servers

You might need to reassociate virtual machine hosts and VMM library servers with the VMM management server after the upgrade.

1. Click **Fabric > Servers > All Hosts**.
2. In the **Hosts** pane, ensure that the **Agent Status** column is displayed. If it isn't right-click a column heading > **Agent Status**.
3. Select the host that you need to reassociate with the VMM management server.
4. In the host group, click **Refresh**. If a host needs to be reassociated, the **Host Status** column will display a value of **Needs Attention**, and the **Agent Status** column will display a value of **Access Denied**. Right-click the host that you want to reassociate, and then click **Reassociate**. In **Reassociate Agent** provide credentials, and then click **OK**. The Agent Status column will display a value of Reassociating. After the host has been reassociated successfully, the **Agent Status** column will display a value of **Responding**. And after you refresh the host again, the **Host Status** column for the host will display a value of **OK**. After you have reassociated the host, you will most likely have to update the VMM agent on the host.

Redeploy Azure Site Recovery

If Azure Site Recovery was integrated into your VMM 2012 R2 deployment, you need to redeploy it with VMM 2016, for [replication to Azure](#), or [replication to a secondary site](#).

Read this [blog entry](#) for details of Hyper-V host support when running VMM 2016.

Connect to Operations Manager

After the upgrade, reconnect VMM to Operations Manager. Note that you shouldn't install any management pack on VMM 2016 RTM. Update Rollup 1 or later must be installed. If you have installed any management packs on the RTM version, uninstall them before you install Update Rollup 1.

Configure Always On Availability Groups

If you upgraded a database that was configured with Always On Availability Groups, you need to complete a few tasks to ensure that the upgraded database is properly configured with Always On Availability Groups.

1. Add the VMM database to the availability group. You can use Microsoft SQL Server Management Studio to perform this task.
2. On the secondary node computer in the cluster that is running SQL Server, create a new sign-in account. Configure the sign-in name so that it is identical to the VMM service account name. Include user mapping to the VMM database, and configure the database owner credentials.
3. Initiate a failover to the secondary node computer that is running SQL Server, and verify that you can restart the VMM service (scvmmsservice).
4. Repeat the last two steps for every secondary node in the cluster that is running SQL Server.

5. If this is a high availability VMM setup, continue to install other high availability VMM nodes.

Update virtual machine templates

All virtual machine templates that were upgraded need to correctly specify the virtual hard disk that contains the operating system.

1. Click **Library** > **Templates** > **VM Templates**.
2. Right-click the template > **Properties** > **Hardware Configuration**, and check disk settings.

Renew certificates for PXE servers

If you have a PXE server in the VMM fabric, you need to remove it from the fabric, and then add it again. This is in order to renew the PXE server certificate, and avoid certificate errors.

Update driver packages

Driver packages that were previously added to the VMM library must be removed and added again to be correctly discovered.

If you plan to assign custom drivers, the driver files must exist in the library. You can tag the drivers in the library, so that you can later filter them by tag. After files are added, when you configure a physical computer profile, you can specify the driver files. VMM installs the specified drivers when it installs the operating system on a physical computer.

In the physical computer profile, you can select to filter the drivers by tags, or you can select to filter drivers with matching Plug and Play (PnP) IDs on the physical computer. If you select to filter the drivers by tags, VMM determines the drivers to apply by matching the tags that you assign to the drivers in the library to the tags that you assign in the profile. If you select to filter drivers with matching PnP IDs, you don't need to assign custom tags.

1. Locate a driver package that you want to add to the library.
2. In the library share that is located on the library server associated with the group where you want to deploy the physical computers, create a folder to store the drivers, and then copy the driver package to the folder.
3. We strongly recommend that you create a separate folder for each driver package, and that you do not mix resources in the driver folders. If you include other library resources such as .iso images, .vhd files or scripts with an .inf file name extension in the same folder, the VMM library server will not discover those resources. Also, when you delete an .inf driver package from the library, VMM deletes the entire folder where the driver .inf file resides.
4. In the VMM console, open the Library workspace. In the **Library** > **Library Servers**, expand the library server where the share is located, right-click the share, and then click **Refresh**. After the library refreshes, the folder that you created to store the drivers appears.
5. Now assign tags if required. In **Library**, expand the folder that you created to store the drivers in the previous procedure, and then click the folder that contains the driver package.
6. In the **Physical Library Objects**, right-click the driver .inf file, and then click **Properties**.
7. In the **Driver File Name Properties** > **Custom tags**, enter custom tags separated by a semi-colon, or click **Select** to assign available tags or to create and assign new ones. If you click **Select**, and then click **New Tag**, you can change the name of the tag after you click **OK**. For example, if you added a network adapter driver file, you could create a tag that is named ServerModel NetworkAdapterModel, where ServerModel is the server model and NetworkAdapterModel is the network adapter model.

Relocate the VMM library

- If you upgraded to a high availability VMM management server, we recommend that you relocate your VMM library to a high availability file server.

- After you create a new VMM library, you will want to move the resources from the previous VMM library to the new VMM library.
- To preserve the custom fields and properties of saved virtual machines in the previous VMM library, deploy the saved virtual machines to a host and then save the virtual machines to the new VMM library.
- Note that operating system and hardware profiles cannot be moved. You need to re-create these profiles.

Next steps

Learn about deploying the latest [update rollups](#).

Install VMM

8/13/2019 • 11 minutes to read

Use this article to install the System Center - Virtual Machine Manager (VMM) management server.

Before you start

- Review the system requirements and [planning information](#). Learn about [system requirements](#)
- Make sure you have at least local admin permissions on the computer before you run the setup.
- The service account should be an administrator on the VMM server.

Run setup

NOTE

The service account for VMM can be:

- A local account
- A user account used for service
- A group managed service account.
- If you are using a local account, you can't have VMM in highly available configuration.
- If you are using gMSA account, the format should be *domainFQDN\gMSAAccount\$*

1. Close any open programs and ensure that no restarts are pending on the computer.
2. To start the Virtual Machine Manager Setup wizard, on your installation media, right-click **setup.exe**, and then click **Run as administrator**.
3. In the main setup page, click **Install**.
4. On the **Select features to install** page, select the **VMM management server** check box, and then click **Next**. The VMM console will be automatically installed. If you're installing on a cluster node you'll be asked if you want to make the management server highly available.
5. On the **Product registration information** page, provide the appropriate information, and then click **Next**. If you do not enter a product key, VMM will be installed as an evaluation version that expires in 180 days after installation.
6. On the **Please read this license agreement** page, review the license agreement, select the **I have read, understood, and agree with the terms of the license agreement** check box, and then click **Next**.
7. On the **Diagnostic and Usage Data** page, review Microsoft's data collection policy and how to disable data collection. Then click **Next**.
8. If the **Microsoft Update** page appears, select whether you want to use Microsoft Update, and then click **Next**. If you've already chosen to use Microsoft Update on this computer the page won't appear.
9. On the **Diagnostic and Usage Data** page, review Microsoft's data collection policy and how to disable data collection. Then click **Next**.
10. On the **Installation location** page, use the default path or type a different installation path for the VMM program files, and then click **Next**. The setup program checks the computer on which you are installing the VMM management server to ensure that the computer meets the appropriate hardware and software requirements. If the computer does not meet a prerequisite, a page that contains information about the prerequisite and how to resolve the issue appears.
11. On the **Database configuration** page, if you're using a remote SQL instance specify the name of the

computer that is running SQL Server. If you are installing the VMM management server on the same computer that is running SQL Server, then in the **Server name** box, either type the name of the computer (for example, **vmmserver01**) or type **localhost**. If the SQL Server is in a cluster, type the cluster name.

12. Don't specify a **Port** value if you don't have a remote instance of SQL Server, or if you have a remote SQL Server that uses the default port (1443).
13. Specify the SQL Server instance name and whether to use an existing or new database. You'll need an account with permissions to connect to the instance.
14. On the **Configure service account and distributed key management** page, specify the account that the VMM service will use. You can't change the identity of the VMM service account after installation. Learn more about distributed key management [here](#).
15. Under **Distributed Key Management**, select whether to store encryption keys in Active Directory.
16. On the **Port configuration** page, use the default port number for each feature or provide a unique port number that is appropriate in your environment. You cannot change the ports that you assign during the installation of a VMM management server unless you uninstall and then reinstall the VMM management server. Also, do not configure any feature to use port 5986, because that port number is preassigned.
17. On the **Library configuration** page, select whether to create a new library share or to use an existing library share on the computer. The default library share that VMM creates is named **MSSCVMLibrary**, and the folder is located at **%SYSTEMDRIVE%\ProgramData\Virtual Machine Manager Library Files**. **ProgramData** is a hidden folder, and you cannot remove it. After the VMM management server is installed, you can add library shares and library servers by using the VMM console or by using the VMM command shell.
18. On the **Installation summary** page, review your selections and then click **Install**. The **Installing features** page appears and displays the installation progress.
19. On the **Setup completed successfully** page, click **Close** to finish the installation. To open the VMM console, you can ensure that **Open the VMM console when this wizard closes** is checked, or you can click the **Virtual Machine Manager Console** icon on the desktop.

During Setup, VMM enables the following firewall rules. These rules remain in effect even if you later uninstall VMM.

- Windows Remote Management
- Windows Standards-Based Storage Management

NOTE

If Setup does not finish successfully, consult the log files in the **%SYSTEMDRIVE%\ProgramData\VMMLogs** folder. **ProgramData** is a hidden folder.

Install VMM from a command prompt

You can install VMM from a command prompt. The installation media contains .ini files for all VMM features:

- **VMServer.ini**: Settings for the VMM management server.
- **VMClient.ini**: Settings for the VMM console.
- **VMServerUninstall.ini**: Uninstallation settings for the VMM management server.

Each of these files contain key/value pairs with default values. These entries are commented out. Remove the comment symbol(#) and change the value.

1. Edit the **VMServer.ini** file with the options in the table below this procedure
2. After you edit open an elevated command prompt and run **setup.exe** with the parameters below. For example, to use a **VMServer.ini** file that is stored in **C:\Temp** with a SQL Server administrator account of

contoso\SQLAdmin01 and a VMM service account of contoso\VMMAdmin14, use the following command:

```
setup.exe /server /i /f C:\Temp\VMServer.ini /SqlDBAdminDomain contoso /SqlDBAdminName SQLAdmin01 /SqlDBAdminPassword password123 /VmmServiceDomain contoso /VmmServiceUserName VMMAdmin14 /VmmServiceUserPassword password456 /IACCEPTSCEULA
```

VMServer.ini values

OPTION	VALUES	DEFAULT
ProductKey	Product key in the format: xxxxx-xxxxx-xxxxx-xxxxx-xxxxx	xxxxx-xxxxx-xxxxx-xxxxx-xxxxx
UserName	Optional display name for the user who is installing the features. This is not the user account for the installation.	Administrator
CompanyName	Optional display name for the organization that is installing the features.	Microsoft Corporation
ProgramFiles	Location for VMM files.	C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager
CreateNewSqlDatabase	0: Use an existing Microsoft SQL Server database. 1: Create a new SQL Server database.	1
SqlInstanceName	Name of the new or existing instance of SQL Server.	MICROSOFT\$VMM\$
SqlDatabaseName	Name of the new or existing SQL Server database.	VirtualManagerDB
RemoteDatabaseImpersonation	0: Do not impersonate the administrator account for SQL Server. The user that runs setup.exe must be an administrator for the server that is hosting SQL Server. 1: Impersonate the administrator account for SQL Server by using the provided credentials. The user who runs setup.exe must provide values for the SqlDBAdminName, SqlDBAdminPassword, and SqlDBAdminDomain parameters.	0
SqlMachineName	Name of the server that is hosting SQL Server. Do not specify localhost. Instead, specify the actual name of the computer.	<sqlmachinename>

OPTION	VALUES	DEFAULT
(various ports)	Ports used by VMM	IndigoTcpPort: 8100 IndigoHTTPSPort: 8101 IndigoNETTCPPort: 8102 IndigoHTTPPort: 8103 WSManTcpPort: 5985 BitsTcpPort: 443
CreateNewLibraryShare	0: Use an existing library share. 1: Create a new library share.	1
LibraryShareName	Name of the file share to be used or created.	MSSCVMMLibrary
LibrarySharePath	Location of the existing file share or the new file share to be created.	C:\ProgramData\Virtual Machine Manager Library Files
LibraryShareDescription	Description of the share.	Virtual Machine Manager Library Files
SQMOptIn	0: Do not opt in for "Diagnostic and Usage Data". 1: Opt in for "Diagnostic and Usage Data"	1
MUOptIn	0: Do not opt in to Microsoft Update. 1: Opt in to Microsoft Update.	0
VmmServiceLocalAccount	0: Use a domain account for the VMM service (scvmm-service). 1: Use the Local System account for the VMM service. To use a domain account, when you run setup.exe, provide values for the VMMServiceDomain, VMMServiceUserName, and VMMServiceUserPassword parameters.	0
TopContainerName	Container for Distributed Key Management (DKM); for example, "CN=DKM,DC=contoso,DC=com".	VMMServer
HighlyAvailable	0: Do not install as highly available. 1: Install as highly available.	0

OPTION	VALUES	DEFAULT
VmmServerName	Clustered service name for a highly available VMM management server. Do not enter the name of the failover cluster or the name of the computer on which the highly available VMM management server is installed.	<VMMServerName>
VMMStaticIPAddress	IP address for the clustered service name for a highly available VMM management server, if you are not using Dynamic Host Configuration Protocol (DHCP). Both IPv4 and IPv6 are supported.	<comma-separated-ip-for-HAVMM>
Upgrade	0: Do not upgrade from a previous version of VMM. 1: Upgrade from a previous version.	1

Setup-exe parameters

PARAMETER	DETAILS
/server	Specifies installation of the VMM management server.
/i or /x	Specifies whether to install (/i) or uninstall (/x) the server.
/f <filename>	Specifies the .ini file to use. Be sure that this parameter points to the correct .ini file. If setup.exe does not find an .ini file, it will perform the installation by using its own default values.
/VmmServiceDomain <domainName>	Specifies the domain name for the account that is running the VMM service (scvmmsservice). Use this parameter only if you set VmmServiceLocalAccount to 0 in VMServer.ini.
/VmmServiceUserName <userName>	Specifies the user name for the account that is running the VMM service (scvmmsservice). Use this parameter only if you set VmmServiceLocalAccount to 0 in VMServer.ini.
/VmmServiceUserPassword <password>	Specifies the password for the account that is running the VMM service (scvmmsservice). Use this parameter only if you set VmmServiceLocalAccount to 0 in VMServer.ini.
/SqlDBAdminDomain <domainName>	Specifies the domain name for the administrator account for the SQL Server database. Use this parameter if the current user does not have administrative rights to SQL Server.
/SqlDBAdminName <userName>	Specifies the user name for the administrator account for the SQL Server database. Use this parameter if the current user does not have administrative rights to SQL Server.
/SqlDBAdminPassword <password>	Specifies the password for the administrator account for the SQL Server database. Use this parameter if the current user does not have administrative rights to SQL Server.

PARAMETER	DETAILS
/IACCEPTSCEULA	<p>Notes acceptance of the Microsoft Software License Terms. This is a mandatory parameter.</p> <p>For example, to use a VMServer.ini file that is stored in C:\Temp with a SQL Server administrator account of contoso\SQLAdmin01 and a VMM service account of contoso\VMAdmin14, use the following command:</p> <pre>setup.exe /server /i /f C:\Temp\VMServer.ini /SqlDBAdminDomain contoso /SqlDBAdminName SQLAdmin01 /SqlDBAdminPassword password123 /VmmServiceDomain contoso /VmmServiceUserName VMAdmin14 /VmmServiceUserPassword password456 /IACCEPTSCEULA</pre>

Uninstall VMM or the VMM console

1. Make sure the VMM console and VMM command shell are closed.
2. On the computer on which the VMM management server is installed, click **Start**, and then click **Control Panel**.
3. Under **Programs**, click **Uninstall a program**. Under **Name**, double-click **Microsoft System Center Virtual Machine Manager**.
4. On the **What would you like to do?** page, click **Remove features**.
5. On the **Select features to remove** page, select the **VMM management server** check box, and then click **Next**. If you want to uninstall the VMM console, select the **VMM console** check box. Note that if you have a highly available VMM deploy, you must remove both the VMM server and VMM console.
6. On the **Database options** page, select whether you want to retain or remove the VMM database, and, if necessary, credentials for the database, and then click **Next**.
7. On the **Summary** page, review your selections and click **Uninstall**. The **Uninstalling features** page appears and uninstallation progress is displayed.
8. After the VMM management server is uninstalled, on the **The selected features were removed successfully** page, click **Close**.

The following firewall rules, which were enabled during VMM Setup, remain in effect after you uninstall VMM:

- File Server Remote Management
- Windows Standards-Based Storage Management firewall rules

If there is a problem with setup completing successfully, consult the log files in the %SYSTEMDRIVE%\ProgramData\VMMLogs folder. **ProgramData** is a hidden folder.

Uninstall VMM from the command line

To uninstall VMM edit the VMServerUninstall.ini file as described below. Then run setup.exe for the uninstall. For example, to uninstall using an ini file stored in C:\Temp with an account contoso.SQLAdmin01 type: **setup.exe /server /x /f C:\Temp\VMServerUninstall.ini /SqlDBAdminDomain contoso /SqlDBAdminName SQLAdmin01 /SqlDBAdminPassword password123**

VMServerUninstall.ini

OPTION	DETAILS	DEFAULT VALUE
--------	---------	---------------

OPTION	DETAILS	DEFAULT VALUE
RemoteDatabaseImpersonation	<p>0: Local SQL Server installation.</p> <p>1: Remote SQL Server installation.</p> <p>When you run setup.exe, provide a value for the SqlDBAdminName, SqlDBAdminPassword, and SqlDBAdminDomain parameters unless the user who is running setup.exe is an administrator for SQL Server.</p>	0
RetainSqlDatabase	<p>0: Remove the SQL Server database.</p> <p>1: Do not remove the SQL Server database</p> <p>To remove the SQL Server database, when you run setup.exe, provide a value for the SqlDBAdminName, SqlDBAdminPassword, and SqlDBAdminDomain parameters unless the user who is running Setup is an administrator for SQL Server.</p>	0
ForceHAVMMUninstall	<p>0: Do not force uninstallation if setup.exe cannot verify whether this node is the final node of the highly available installation.</p> <p>1: Force the uninstallation.</p>	

Support for gMSA account

Group Managed Service Account (gMSA) helps improve the security posture and provides convenience through automatic password management, simplified service principle name (SPN) management and the ability to delegate the management to other administrators.

VMM 2019 supports the use of gMSA for *Management server service account*.

Prerequisites

1. Review [this article](#) and create gMSA as per the guidance available in the article.
2. Make sure that the servers on which VMM Management service would be installed have permissions to retrieve the password of gMSA account.

NOTE

You do not need to specify the SPN when creating the gMSA. VMM service sets the appropriate SPN on the gMSA.

Use the following steps

1. Start the VMM installation setup.
2. On the **Service account configuration** page, select **Group Managed Service Account** as the option for VMM service account.
3. Enter the gMSA account details in *Domain\gMSA account* format.

Microsoft System Center Virtual Machine Manager Setup Wizard

×

Configuration

Report a problem

Configure service account and distributed key management

Virtual Machine Manager Service Account

Select the account to be used by the VMM service. Highly available VMM installations require the use of a domain account.
[Which type of account should I use?](#)

☐ Local System account

☐ Domain account

☒ Group Managed Service Account

User name and domain:

Password:

Select...

User name and domain:

contoso\demogMSA

Distributed Key Management

Select whether to store encryption keys in Active Directory instead of on the local machine. Highly available VMM installations require the keys be stored in Active Directory.

☐ Store my keys in Active Directory

Provide the location in Active Directory. For example, CN=DKM,DC=contoso,DC=com.

[How do I configure distributed key management?](#)

Previous

Next

Cancel

Install the VMM Console

6/14/2019 • 3 minutes to read

This article describes how to install the System Center Virtual Machine Manager (VMM) console on a remote computer, and connect to the VMM server. When you install the VMM management server, the console is installed on it automatically.

Before you start

- Check the supported operating systems for the console that are detailed in the system requirements.
- Review and ensure you meet the system requirements. [Learn about](#) system requirements.
- Make sure you have at least local administrator permissions on the computer on which you're installing the console.
- The VMM version of the console must match the System Center version of the VMM server. For example, to connect to a VMM server running System Center 2016, the VMM console version must also be 2016.
- You can only install one version of the console on a single machine.

Run setup

1. Review the planning instructions. Then, right-click setup.exe for VMM > **Run as administrator**.
2. On the main setup page, click **Install** and on the **Select features to install** page, select the **VMM console** check box, and then click **Next**. On the **Please read this notice** page, select the **I agree with the terms of this notice** check box, and then click **Next**.
3. Review the information on the **Diagnostic and Usage Data** page, and then click **Next**. On the **Microsoft Update** page, select whether you want to use Microsoft Update, and then click **Next**. This page won't appear if updates are already installed.
4. On the **Installation location** page, type an installation path for the VMM program files or use the default path, and then click **Next**. Setup checks that the computer meets the console installation requirements.
5. On the **Port configuration** page, type the port that you want to use for the VMM console to communicate with the VMM management server, and then click **Next**. The port setting that you assign for the VMM console should match the port setting that you assigned for the VMM console during the installation of the VMM management server. The default port setting is 8100. Also, do not assign port number 5986, because it is preassigned.
6. On the **Installation summary** page, review the settings and click **Install**. The **Installing features** page appears and displays the installation progress.
7. On the **Setup completed successfully** page, click **Close** to finish the installation. Select **Open the VMM console when this wizard closes** to open the console after the wizard finishes. If setup does not finish successfully, consult the log files in the %SYSTEMDRIVE%\ProgramData\VMMLogs folder. **ProgramData** is a hidden folder.

If you want to uninstall the console do the following:

1. In the Control Panel click to uninstall a program and select Microsoft System Center Virtual Machine Manager.
2. In **Select features to remove** click **VMM console**. On the VMM management server, you can't uninstall the console without uninstalling VMM management.
3. In **Summary** review the settings and click **Uninstall**.

Install the console from the command prompt

You can install the VMM administrator console from the command line by using the VMClient.ini file to customize the installation options. You can set the parameters as follows:

PARAMETER	VALUE
ProgramFiles	Specify the location in which to store program files.
IndigoTCpPort	Specify the port number used to communicate with the VMM server.
MUOptIn	0: Don't opt in to Microsoft Update. 1: Opt in.
VmmServerForOpsMgrConfig	Specify the name of the System Center Operations Manager server.

Install the console as follows:

1. Copy the VMClient.ini file from the amd64\Setup folder to a local folder.
2. Edit the VMClient.ini file. Remove the comment indicator (#) only if you want to edit the entry. Otherwise, setup uses the values in the .ini file as the default values.
3. Run setup as follows: `setup.exe /client /i /f <path>`, where:
 - `/client` - specifies console installation
 - `/i` or `/x` - specify whether to install (`/i`) or uninstall (`/x`) the console.
 - `/f <filename>` - specifies the ini file to use. Make sure this is correct. If setup doesn't find the ini file, it installs with default values.
 - `path`: location of ini file.
 - Don't use the `/opsmgr` parameter

Example: **`setup.exe /client /i /f C:\Temp\VMClient.ini`**.

Connect to a VMM management server

1. On the remote machine desktop click the VMM console icon if the console isn't open.
2. Specify the server name and port on which the VMM server is listening (8100 by default). For example: **`vmmserver01:8100`**.
3. You can use the account with which you're logged on to the remote machine or specify a different account to connect to VMM. Click **Automatically connect with these settings** if you want to connect to the same VMM server each time.
4. Click **Connect**. **Select User Role** appears if the credentials you're using belong to more than one VMM user role.

Enable enhanced console session in VMM

5/30/2019 • 2 minutes to read

This article provides information about how to configure enhanced console session in System Center - Virtual Machine Manager (VMM).

Console connect in VMM provides a way to connect to the VM without a network connection to it. For information on deploying VMM console, see [install VMM console](#). Prior to 1801, the console connect in VMM supported only basic session where clipboard text can only be pasted through **Type Clipboard Text menu** option.

VMM 1801 and later supports enhanced console session. With console connected through enhanced session, **Cut (Ctrl + X)**, **Copy (Ctrl + C)** and **Paste (Ctrl + V)** operations on the ANSI text and files are available on the clipboard, thereby copy/paste commands for text and files are made possible from and to the VM.

Before you start

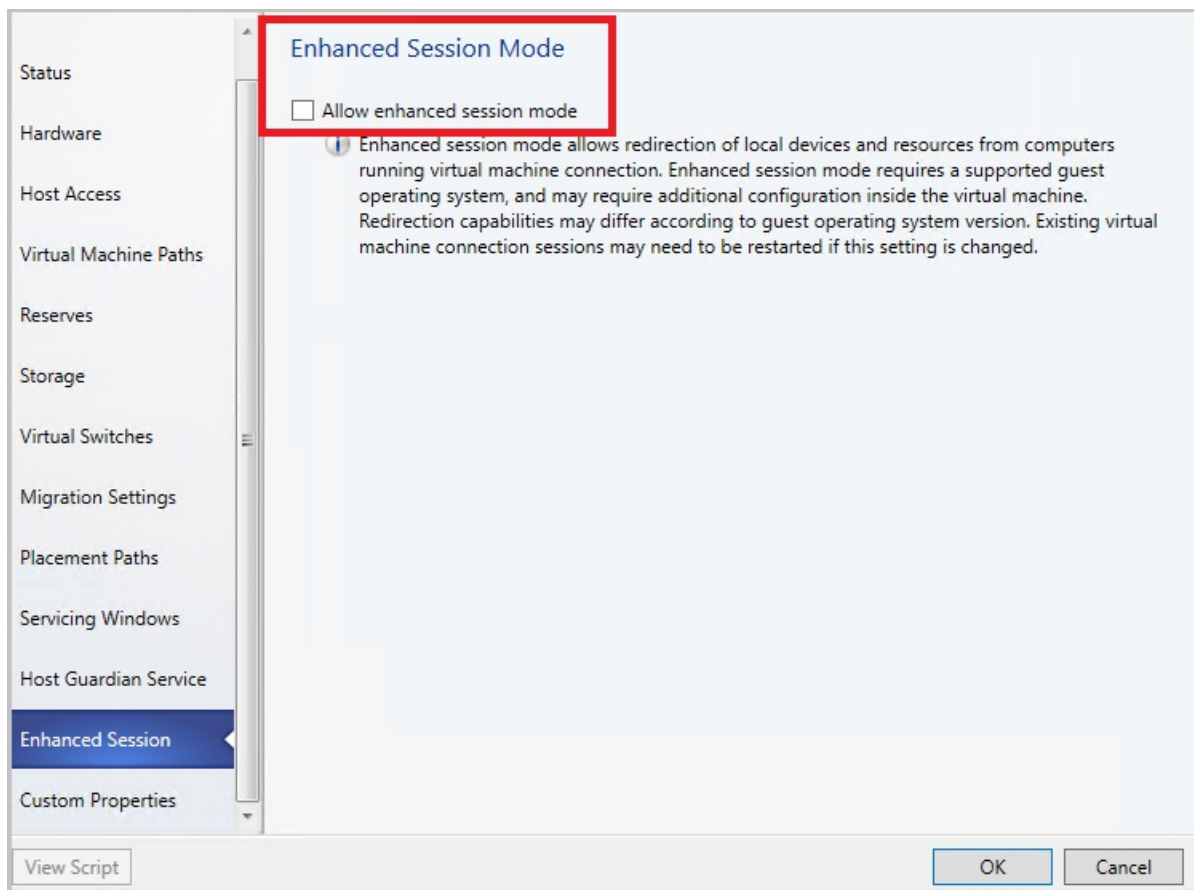
Ensure the following prerequisites:

1. The operating system of the host, on which the VM is running, should be Windows Server 2016 and above.
2. The Hyper-V host must have Enhanced session mode policy setting turned ON.
3. The computer from which you connect to the VM must run on Windows 10, Windows 8.1, Windows Server, Windows Server 2016 or above.
4. The virtual machine must have remote desktop services enabled and run Windows 10, Windows 8.1, Windows Server 2016 or Windows Server 2012 R2 as the guest operating system.


Procedure - enable the enhanced console session

Use the following steps:

1. Right-click the host in **VMs & Services** and navigate to **Enhanced Session** option.
2. Select **Allow enhanced session mode** and click **OK**.



3. In VMM console, navigate to the VM on this host.
4. Right-click the VM and click **Connect via Console**.
5. Once you see the VM console, look for **Open Enhanced Session** option at the top right of the page. Click it to launch the enhanced Session window.



NOTE

This action will close the current session and opens a new session. You will be redirected to the login screen in the new session.

6. To switch back to basic session, click **Open Basic Session** at the top right.

NOTE

Once Enhanced Session Mode policy is enabled on the host:

- Close any open console sessions to view the **Enhanced Session** option.
- For a VM that is booting for the first-time from a VHD/VHDX, the enhanced session option doesn't show up when you attempt to connect through console. Restart the VM and refresh the VM properties in VMM for the Enhanced Session option to appear in console connect window.

Next steps

[Use local resources on Hyper-v VM with VMconnect](#)

Deploy VMM for high availability

6/6/2019 • 2 minutes to read

For resilience and scalability you can deploy System Center - Virtual Machine Manager (VMM) in high availability mode.

Before you start

Prepare for a high availability deployment by considering the following:

- Only one instance of VMM can be deployed to a failover cluster of up to 16 nodes.
- Requirements for computers running as VMM management nodes:
 - All cluster nodes that are VMM servers must be running Windows Server 2016.
 - Each cluster node must be joined to a domain and must have a computer name that does not exceed 15 characters.
 - The VMM service network name must not exceed 15 characters.
 - Windows ADK needs to be installed on each computer. Install from setup or the [download center](#). Select **Deployment Tools** and **Windows Preinstallation Environment** when you install.
 - If you plan to deploy VMM services that use SQL Server data-tier applications, install the related command-line utilities on your VMM management server. The command line utility is available in the [SQL Server 2012 feature pack](#) or [SQL Server 2014 feature pack](#) or [SQL Server 2016 feature pack](#).
- Don't install on a Hyper-V host parent partition. You can install VMM on a VM.
- Before you start, set up the VMM service account and distributed key management. [Learn more](#)

Deploy high availability components

- [Deploy the VMM management server in a failover cluster](#)
- [Make library server file shares highly available](#)
- [Deploy the SQL Server VMM database as highly available](#)

Deploy a highly available VMM management server

12/7/2018 • 3 minutes to read

This article describes the steps for deploying a highly available System Center - Virtual Machine Manager (VMM) server.

Before you start

- Read the [planning steps](#) for a highly available deployment.
- This procedure presumes you're setting up a single failover cluster with two or more file servers.

Set up the failover cluster

1. Click **Server Manager** > **Manage** > **Add Roles and Features**.
2. In **Select installation type** click **Role-based or feature-based installation**.
3. In **Select destination server** click the server you want to configure for failover clustering. On **Select features** click **Failover Clustering**. Click **Add Feature** to install the failover cluster management tools.
4. In **Confirm installation selections** click **Install**. A server restart isn't needed.
5. Repeat for each server you want to add as a node in the file server cluster.
6. After you have at least two nodes in the cluster you can run cluster validation tests (Open **Failover Cluster Manager** and under **Management** click **Validate Configuration**).
7. In **Select Servers or a Cluster** specify the NetBIOS or FQDN of a node you're adding and click **Add**. In **Testing Options** click **Run all tests (recommended)**.
8. In **Summary**, if the tests completed correctly click **Create the cluster now using the validated nodes**. Click **View Report** to troubleshoot any issues.
9. In **Access Point for Administering the Cluster** specify the cluster name. For example **VMMLibrary**. When the cluster is created this name is registered as the cluster computer object (CNO) in Active Directory. If you specify a NetBIOS name for the cluster the CNO is created in the same location where the computer objects for the cluster node reside (either the default Computers container or an OU). You can specify a different location by adding the distinguished OU name. For example CN=ClusterName, OU=Clusters,DC=Contoso. For more information, see [distributed key management](#).
10. If the server isn't configured to use DHCP specify a static IP address for the cluster. Select each network you want to use for cluster management and in **Address** select the IP address. This is the IP address that is associated with the cluster in DNS.
11. In **Confirmation** review the settings. Clear **Add all eligible storage to the cluster** if you want to configure storage later. Click **Next** to create the cluster.
12. In **Summary** confirm that the cluster was created and that the cluster name is listed in Failover Cluster Manager.

Install VMM on the first cluster node

1. On either node of the cluster you created run VMM setup and click **Install**.
2. VMM detects its installing on a cluster node and asks you if you want to make the VMM server highly available. Click **Yes**.
3. In **Select features to install** select the VMM management server and the VMM console.
4. In **Product registration information** specify organizational details and the product key.
5. IN **EULA and CEIP** accept the EULA and specify whether you want to participate in CEIP.
6. In **Installation Location** accept the default settings.

7. In **Prerequisites** VMM checks whether all prerequisites are met and installs missing components. If you don't have the Windows ADK installed, download and install it.
8. In **Database configuration** specify the database to use for VMM. The database should be highly available and deployed in a separate failover cluster. This dialog appears if VMM isn't clustered, or if it's clustered but not using Always On Availability Groups. Specify the cluster name.
9. In **Cluster configuration** specify the name of the VMM cluster for example **HAVMMM**.
10. In **Configure service account and distributed key management** specify the service account and key location you created earlier. VMM Run As accounts are stored as encrypted in the VMM database. For a high availability deployment, you need to access encrypted keys from a central location so you should have created a distributed key management container in Active Directory before you ran setup. Learn more about distributed key management container [here](#).
11. In **Port configuration** modify port settings if you need to.
12. Finish installing VMM. You can't specify a library share right now. In a highly available deployment you create the library share after installation is complete.

Install VMM on the second cluster node

1. Run setup and confirm that you want to **add this server as a node** to the highly available deployment.
2. During the wizard specify the service account password. You don't need to specify other information.

Deploy SQL Server for VMM high availability

8/22/2019 • 2 minutes to read

This article describes the steps for deploying a highly available SQL Server database for System Center - Virtual Machine Manager (VMM). You set up a SQL Server cluster, configure the SQL Server VMM database with Always On Availability Groups.

Before you start

Read the [planning information](#) for a highly available VMM deployment. It includes prerequisites and issues you should be aware of.

Set up availability groups

SQL Server Always On availability groups support failover environments for a discrete set of user databases (availability databases). Each set of availability databases is hosted by an availability replica. To set up an availability group you must deploy a Windows Server Failover Clustering (WSFC) cluster to host the availability replica, and enable Always On on the cluster nodes. You can then add the VMM SQL Server database as an availability database.

- [Learn more](#) about Always On prerequisites
- [Learn more](#) about setting up a WSFC for Always On availability groups
- [Learn more](#) about setting up an availability group

Configure the VMM database with Always On Availability Groups

1. On the VMM server, stop the VMM service. For a cluster, in Failover Cluster Manager, stop the VMM role.
2. Connect to the machine that hosts the VMM database, and in SQL Server Management Studio, right-click the VMM database > **Properties**. In **Options**, set the **Recovery model** for the database to **Full**.
3. Right-click the VMM database > **Tasks** > **Back Up** and take a backup of the database.
4. In SQL Server Management Studio > **Always On High Availability** > right-click the availability group name > **Add Database**.
5. In **Add Database to Availability Group** > **Select Databases**, select the VMM database.
6. In **Select Data Synchronization**, leave the **Full** default.
7. In **Connect to Replicas** > **Connect**, specify permissions for the availability group destination.
8. Prerequisites are checked in **Validation**. In **Summary**, when you click **Next** Always On availability support is initiated for the VMM database. The VMM database is copied and from this point Always On keeps the VMM database synchronized between the SQL Server Always On cluster nodes.
9. Restart the VMM service or cluster role. The VMM server should be able to connect to the SQL Server.
10. VMM credentials are only stored for the main SQL Server, so you need to create a new login on the secondary node of the SQL Server cluster, with the following characteristics:
 - The login name is identical to the VMM service account name.
 - The login has the user mapping to the VMM database.
 - The login is configured with the database owner credentials.

Run a failover

To check that Always On is working as expected for the VMM database, run a failover from the primary to

secondary node in the SQL Server cluster.

1. In SQL Server Management Studio, right-click the availability group on the secondary server > **Failover**.
2. In **Fail Over Availability Group** > **Select New Primary Replica**, select the secondary server.
3. In **Summary**, click **Finish**.
4. Now move it back by initiating a failover to the secondary node computer that is running SQL Server, and verify that you can restart the VMM service (scvmmsservice).
5. Repeat the last two steps for every secondary node in the cluster that is running SQL Server.
6. If this is a high availability VMM setup, continue to install other high availability VMM nodes.

NOTE

If you are experiencing high latency or timeout errors in a multi-subnet scenario, change VMM connection string in the path *HKLM\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings\Sql\ConnectionString*, add `MultiSubnetFailover=True`; and restart the VMM service.

Deploy a highly available VMM library

12/7/2018 • 3 minutes to read

This article describes the steps for deploying a highly available System Center - Virtual Machine Manager (VMM) library. You set up a Windows failover cluster running the File Server role. Then you create file shares on the cluster, and assign them as VMM library shares.

Before you start

Read the [planning steps](#) for a highly available VMM deployment.

Set up the failover cluster

This procedure presumes you're setting up a single failover cluster with two or more file servers.

1. Click **Server Manager > Manage > Add Roles and Features**.
2. In **Select installation type** click **Role-based or feature-based installation**.
3. In **Select destination server** click the server you want to configure for failover clustering. On **Select features** click **Failover Clustering**. Click **Add Feature** to install the failover cluster management tools.
4. In **Confirm installation selections** click **Install**. A server restart isn't needed.
5. Repeat for each server you want to add as a node in the file server cluster.
6. After you have at least two nodes in the cluster you can run cluster validation tests (you'll need at least two nodes in the cluster). Open **Failover Cluster Manager** and under **Management** click **Validate Configuration**.
7. In **Select Servers or a Cluster** specify the NetBIOS or FQDN of a node you're adding and click **Add**. In **Testing Options** click **Run all tests (recommended)**.
8. In **Summary**, if the tests completed correctly click **Create the cluster now using the validated nodes**. Click **View Report** to troubleshoot any issues.
9. In **Access Point for Administering the Cluster** specify the cluster name. For example **VMMLibrary**. When the cluster is created this name will be registered as the cluster computer object (CNO) in Active Directory. If you specify a NetBIOS name for the cluster the CNO is created in the same location where the computer objects for the cluster node reside (either the default Computers container or an OU). You can specify a different location by adding the distinguished OU name. For example CN=ClusterName, OU=Clusters,DC=Contoso.
10. If the server isn't configured to use DHCP specify a static IP address for the cluster. Select each network you want to use for cluster management and in **Address** select the IP address. This is the IP address that will be associated with the cluster in DNS.
11. In **Confirmation** review the settings. Clear **Add all eligible storage to the cluster** if you want to configure storage later. Click **Next** to create the cluster.
12. In **Summary** confirm that the cluster was created and that the cluster name is listed in Failover Cluster Manager.

If you want to build a guest cluster to deploy the file server, read Rudolf Vesely's useful [blog post](#).

Set up the file server role

1. On each computer you'll set up as a file server node, in **Failover Cluster Manager** select **Configure Role**.
2. In the **High Availability Wizard > Select Role**, select **File Server**.

3. In **File Server Type** select **File Server for general use**.
4. In **Client Access Point** type the cluster name (in our procedure this was **VMMLibrary** and the cluster IP address).
5. In **Select storage** specify the shared storage you want to use.
6. Confirm the settings and finish the wizard.

Create a file share

1. In **Failover Cluster Management** > cluster name > **Roles**. Select the file server and click **Add File Share**.
2. In **New Share Wizard** > **Select Profile** select **SMB Share - Quick**.
3. In **Share Location** select the file server.
4. In **Share Name** specify the share name and description.
5. In **Other Settings** leave the default settings.
6. In **Permissions** grant full access to the SYSTEM and Administrators accounts, and to the VMM admin account. In **Confirmation** review settings and click **Create**.

Add the share as a VMM library

1. Open the VMM console > **Library** > **Add Library Server**.
2. In **Add Library Server Wizard** > **Enter Credentials** specify a domain account with permissions for the file cluster.
3. In the **Select Library Servers** page, type in the domain in which the file cluster is located and in **Computer name** specify the name you assigned to file server cluster or click **Search** to find it. Click **Add** > **Next**.
4. On the **Add Library Servers** page select the library shares you want to add. If you want to add the default library resources to the share select **Add Default Resources**. In addition to default resources this adds the ApplicationFrameworks folder to the share.
5. On the **Summary** page review settings and click **Add Library Servers**. In **Library** > **Library Servers** verify the library server and share are listed.
6. After the share is created you can copy resources to the library share.

Set up TLS for VMM

8/29/2019 • 3 minutes to read

This article describes how to set up Transport Security Layer (TLS) protocol version 1.2 with System Center - Virtual Machine Manager (VMM) server.

Before you start

- VMM should be running VMM 2016 with [Update Rollup 4](#), or later.
- Security fixes should be up-to-date on the VMM server, and the server running the VMM database.
- The VMM server should be running .NET version 4.6. Follow [these instructions](#) to determine which version of .NET is installed.
- To work with TLS 1.2, System Center components generate SHA1 or SHA2 self-signed certificates. If SSL certificates from a certificate authority (CA) certificates are used, they should use SHA1 or SHA2.

Install a SQL Server update for TLS 1.2 support

1. Open [KB 3135244](#).
2. [Download and install](#) the update for your SQL Server version.
 - You don't need this update if you're running SQL Server 2016.
 - SQL Server 2008 R2 doesn't support TLS 1.2.

Configure the VMM server to use TLS 1.2

Disable all SCHANNEL protocols except for TLS 1.2.

Manually modify the registry

1. Open the Registry Editor, and navigate to HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols.
2. Right-click **Protocol**, select **New > Key**. Type in the key and press Enter. Perform this procedure to create the following keys:
 - **SSL3**
 - **TLS 1.0**
 - **TLS 1.1**
 - **TLS 1.2**
3. After you've created these keys, you need to create the **Client** and **Server** keys under them.
 - For **SSL3**, select **New > Key**. Type in **Client** and press enter. Again, for **SSL3**, select **New > Key** again. Then type in **Server** and press Enter.
 - Repeat the action to create the **Client** and **Server** keys under **TLS 1.0**, **TLS 1.1**, and **TLS 1.2**.
4. After you've created the **Client** and **Server** keys, you need to create DWORD values under them, in order to enable and disable protocols. Do this as follows:
 - Enable the TLS 1.2 protocol. To do this, in **TLS 1.2**, under the **Client** key, create the DWORD value **DisabledByDefault**, and set the value to 0. Now create a DWORD value **Enabled**, and set the value to 1. Create the same DWORD values under the **Server** key.
 - Now disable the other protocols. To do this, in **SSL3**, **TLS 1.0** and **TLS 1.1**, under the **Client** key, create the DWORD value **DisabledByDefault**, and set the value to 1. Now create a DWORD value **Enabled**, and set the value to 0. Create the same DWORD values under the **Server** key.

Modify the registry with a PowerShell script

Instead of modifying the registry values manually, you can use the following PowerShell script.

```
$ProtocolList      = @("SSL 2.0","SSL 3.0","TLS 1.0", "TLS 1.1", "TLS 1.2")
$ProtocolSubKeyList = @("Client", "Server")
$DisabledByDefault = "DisabledByDefault"
$Enabled = "Enabled"
$registryPath = "HKLM:\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\"

foreach($Protocol in $ProtocolList)
{
    Write-Host " In 1st For loop"
    foreach($key in $ProtocolSubKeyList)
    {
        $currentRegPath = $registryPath + $Protocol + "\" + $key
        Write-Host " Current Registry Path $currentRegPath"
        if(!(Test-Path $currentRegPath))
        {
            Write-Host "creating the registry"
            New-Item -Path $currentRegPath -Force | out-Null
        }
        if($Protocol -eq "TLS 1.2")
        {
            Write-Host "Working for TLS 1.2"
            New-ItemProperty -Path $currentRegPath -Name $DisabledByDefault -Value "0" -PropertyType
            DWORD -Force | Out-Null
            New-ItemProperty -Path $currentRegPath -Name $Enabled -Value "1" -PropertyType DWORD -Force
            | Out-Null
        }
        else
        {
            Write-Host "Working for other protocol"
            New-ItemProperty -Path $currentRegPath -Name $DisabledByDefault -Value "1" -PropertyType
            DWORD -Force | Out-Null
            New-ItemProperty -Path $currentRegPath -Name $Enabled -Value "0" -PropertyType DWORD -Force
            | Out-Null
        }
    }
}

Exit 0
```

Configure VMM to use TLS 1.2

1. Open the registry editor on the VMM server. Navigate to **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NetFramework\v4.0.30319**.
2. Create the DWORD value **SchUseStrongCrypto**, and set the value to 1.
3. Now navigate to **HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Microsoft\.NetFramework\v4.0.30319**.
4. Under this location, create the same DWORD value **SchUseStrongCrypto**, and set the value to 1.
5. Restart the server for the settings to take effect.

Modify the registry with a PowerShell script

You can modify the registry settings using the following PowerShell script.

```
# Tighten up the .NET Framework
$NetRegistryPath = "HKLM:\SOFTWARE\Microsoft\.NETFramework\v4.0.30319"
New-ItemProperty -Path $NetRegistryPath -Name "SchUseStrongCrypto" -Value "1" -PropertyType DWORD -Force | Out-Null

$NetRegistryPath = "HKLM:\SOFTWARE\WOW6432Node\Microsoft\.NETFramework\v4.0.30319"
New-ItemProperty -Path $NetRegistryPath -Name "SchUseStrongCrypto" -Value "1" -PropertyType DWORD -Force | Out-Null
```

Next steps

[Learn more](#) about the TLS 1.2 protocol.

Deploy and manage update rollups in VMM

12/7/2018 • 5 minutes to read

This article provides information about how to install, verify, and remove update rollups for System Center Virtual Machine Manager (VMM).

NOTE

Update rollups are installed automatically if Microsoft Updates are set to Automatic. If they are not installed automatically, you can install these manually.

Obtain and install update rollups

The following sections provide information about how to obtain and install the update rollups manually through Microsoft updates, Microsoft Download Center and through VMM command line.

Back up the VMM database

1. In the VMM console, click **Settings**.
2. On the Home tab, click **Backup**.
3. In the VMM backup page, in the **Path** text box, specify the location for the backup file, and then click **OK**.

NOTE

Verify that the job completes successfully, otherwise review the error details, and redo the backup.

More information: [Back up the VMM database](#)

Install update rollups from Microsoft updates

Follow these steps on a computer that has a VMM component installed and latest updates are yet to be installed:

1. Under **Settings > Update & Security**, click **Windows Update**.
2. In the **Windows Update** window, click **Check Online for updates from Microsoft Update**.
3. Click **Important updates are available**.
4. Select the update rollup package and then click **OK**.
5. Click **Install updates** to install the update package.

Install update rollups from Microsoft Download Center

To manually download the update rollups from the Microsoft download center, [visit](#) the Microsoft Download Center, search for the desired KB and manually download the update rollup packages.

Install update rollups through command prompt

To manually install an update rollup package through an elevated command prompt, run the following command:

msiexec.exe/update <package_name>

NOTE

In the **<package_name>** placeholder, enter the actual package name.

For example, to install the Update Rollup 2 for System Center 2016 Virtual Machine Manager server (KB3209586), run the following command:

msiexec.exe/update kb3209586_VMMserver_amd64.msp

NOTE

There may be additional installation steps specific to an update rollup release. Check update rollup KB guide to ensure you complete all the installation steps.

Check the installation of an update rollup

Use the following procedure to check if an update rollup is successfully installed:

1. Under **Control Panel > Programs and Features > View installed updates**.
2. Verify that an update entry was created after the update rollup was installed.
 - For example, the Update Rollup 2 was released as update 3209586. You should be able to see this detail under **View installed updates** if you have installed this update rollup and if it was successfully installed.
3. Verify that the binary's version has the correct build number.
4. See the following website to check the build number for a specific update rollup:
 - [List of build numbers for VMM](#)

NOTE

Not all the binaries have the current update rollup build number. However, if you do not have the binaries listed with the relevant update rollup build number, it is likely that the update rollup did not install successfully.

Remove an update rollup

IMPORTANT

1. It is not recommended that you remove the update rollups. Prior to removing, check all the available options if you can avoid the uninstall. You may also contact Microsoft Support to check and ensure if uninstallation is required.
2. It is recommended that you [back up your VMM database](#) before you attempt to remove an update rollup.
3. When you remove, the VMM binaries roll back to their earlier versions. However, the VMM database does not roll back.
4. If you have one or more hotfixes installed on the server, make sure that you replace the hotfix binary with the official update rollup binary, before you start the removal.

Remove an update rollup by using the control panel

1. Go to **Control Panel > Programs > Programs and Features**. For quick access, type **appwiz.cpl** in **Run** it opens the **Programs and Features** in **Control Panel**.
2. Click **View installed updates**.
3. Find the update that you want to remove, right-click the update, and then select **Uninstall**.

Remove an update rollup by using the command line

NOTE

To remove an update by using the command line, you must have the following two globally unique identifiers (GUIDs) available:

- RTM product GUID
- Patch GUID

RTM Product GUIDs

- SCVMM Server: {EBC28D9B-9565-46F3-A248-E26F07F81A98}
- SCVMM Admin Console amd64: {B703D43A-ABF6-4A36-84CC-00D77FF8570B}
- SCVMM Admin Console i386: {F5D46892-E1BD-4E0A-BD6E-DAA1900BA786}
- SCVMM Guest Agent amd64: {3E71E1FB-AF93-4110-A8EB-973132A3B16B}
- SCVMM Guest Agent i386: {57D2C983-23BF-4840-B784-BDDAC2DC932B}

Patch GUIDs

To find the patch GUID, right-click the update, select **Properties**. On the **Details** tab, and then select **Revision number**. This is the patch GUID.

When you know the RTM product GUID and the patch GUID, run the following command to remove the update:
Msiexec /I {< RTM Product GUID >} MSPATCHREMOVE={< Patch GUID >}

Verify the successful removal of an update rollup

Use the following procedure to check if the update rollup was successfully removed:

1. Check whether the update was removed from **Programs and Features**.
2. Verify that binaries were reverted successfully. To do this, go to the VMM installation directory, and verify that there is no binary that has the build version of the update rollup that you uninstalled.

NOTE

When you uninstall the server and console updates from the VMM server, the order of uninstallation is not important.

Restore the database backup in VMM

NOTE

1. You can create a database backup at any time. But the database should be restored only if you uninstalled the latest update rollup on the server. If you uninstalled an older update rollup, you do not have to restore the VMM database.
2. Ensure that you restore the database backup that you created prior to the update rollup installation.

For information on how to restore the database back up, see [this article](#).

Post-restore tasks

This section applies only if you add a host or a library server after you create a database backup and before you uninstall an update rollup.

After the VMM database is recovered, you must do the following actions:

- Remove any hosts that were removed from VMM since the last backup.
- If a host was removed since the last backup, the host shows a status of **Needs Attention** in the VMM console. Also, the virtual machines on that host show a status of **Host Not Responding**.
- Remove the virtual machines that were removed from the VMM since the last back up. If a host has a virtual machine that was removed since the last backup, the virtual machine shows a status of **Missing** in the VMM console.
- Add any hosts that were added to the VMM since the last backup.
- If you restored the VMM database to a different computer, you must re-associate the hosts that have a status of **Access Denied** in the VMM console.
 - A computer is considered different if it has a different security identifier (SID). For example, if you reinstall the operating system on the computer, the computer has a different SID, even if it has the same computer name.

For more information on post-restore tasks, see [this article](#).

NOTE

You must follow these steps for any other fabric resource such as storage providers, library servers, and so on that you add after you create a database backup.

Back up and restore VMM

12/7/2018 • 13 minutes to read

This article describes the backup and recovery process in a System Center - Virtual Machine Manager (VMM) environment, and provides some recommendations.

Before you start

- Don't use checkpoints for disaster recovery. Checkpoints do not create full duplicates of the hard disk contents, nor do they copy data to a separate volume.
- You can use a checkpoint to serve as temporary backup, before updating an operating system on a virtual machine. This allows you to roll back the update if it has adverse effects.
- You should use a backup application to back up and recover your data in case of catastrophic data loss. One option is System Center Data Protection Manager (DPM).
- Data such as Remote Access Authorization (RAA) passwords and the product key can be entered when you reinstall VMM. However, some encrypted data such as Virtual Machine Roles cannot be reentered.
- You can't back up and restore such data if you use the Data Protection application programming interface (DPAPI) for backing up VMM.
- The data will be lost if the VMM management server fails.

Create and implement a backup plan

Basic elements of a backup plan include a list of what needs to be backed up, and an outline of what is changed frequently (and therefore need to be backed up frequently) in your environment.

Back up the VMM database

The VMM database contains information such as configurations, service templates, profiles, virtual machine templates, services, scale-out services, and other critical data that is required for VMM to function correctly. Back up the VMM database regularly.

The VMM database can be stored on the VMM management server or on a separate server running Microsoft SQL Server. To back up the VMM database, you can use one or more of the following:

- SQL Server tools. For more information, see [Create a Full Database Backup \(SQL Server\)](#).
- Other backup tools that are used in your environment.
- The VMM console or a Windows PowerShell cmdlet, as described in the procedures that follow.

In addition to backing up the database, we recommend that you create a system state backup of the VMM management server so that you can re-create the server with the same security identifier (SID) in case of a catastrophic data loss. The SID is an integral part of how VMM is authorized on virtual machine hosts.

IMPORTANT

- There are several ways to recover the VMM database file that you create through either of the following backup procedures. One way, which requires the VMM management server to be functioning, is to use the **SCVMMRecover.exe** tool, as described in [Backup-SCVMMServer](#) (although **SCVMMRecover.exe** is not a cmdlet). Another way, which does not require the VMM management server to be functioning, is to restore by using SQL Server tools for restoring and attaching a database file.
- To use the following procedures, you must be a member of the Administrator user role.

Back up the VMM database by using the VMM console

1. In the **Settings** workspace, on the **Home** tab, in the **Backup** group, click **Backup**.
2. In the **Virtual Machine Manager Backup** dialog box, specify the location for the backup file. Select a folder that is not a root directory and that SQL Server can access.

You can check the status of the backup in the Jobs workspace.

For information about how to recover the backup, see the **Important** note before this procedure.

Back up the VMM database by using cmdlets in Windows PowerShell

1. Start a Windows PowerShell session.
2. At the Windows PowerShell command prompt, run the [Get-SCVMMServer](#) and [Backup-SCVMMServer](#) cmdlets, using the following syntax:

```
get-scvmmserver <VMM management server name> | backup-scvmmserver -Path <BackupFileDir>
```

For information about how to recover the backup, see the **Important** note before these procedures.

Back up hosts and virtual machines

Virtual machine hosts are Hyper-V hosts, VMware ESXi hosts, and host clusters on which virtual machines and services are deployed. To back up virtual machine hosts and clusters, use Microsoft System Center Data Protection Manager (DPM) or another backup application that takes advantage of Volume Shadow Copy Service (VSS) to copy host and virtual machine data to a remote file server share.

IMPORTANT

We recommend that you back up virtual machine configuration files (.vmc) daily.

Inventory your hosts, and then back up all of the hosted virtual machines. To get the list of hosts that are being managed by VMM, run the following cmdlet from a Windows PowerShell command line:

```
$vmhost = get-scvmmserver <VMM management server name> | get-scvmhost
```

For more information, see [Get-SCVMMServer](#) and [Get-SCVMHost](#).

Back up all of the configuration and resource files on each VMM host by using backup software that supports the VMM VSS writer. Backup software that supports VMM can minimize the number of steps required to archive and restore virtual machines, help minimize downtime, and help ensure consistency of the data that is being archived or restored.

Back up library servers

The VMM library includes file-based resources, such as virtual hard disks, ISO images, scripts, driver files, and application packages that are stored on library servers. These resources are closely associated with resources in the VMM database that are not file-based, such as virtual machine and service templates and profiles. All of these resources should be backed up.

To back up the data on library servers, use System Center Data Protection Manager (DPM) or another backup application that takes advantage of Volume Shadow Copy Service (VSS) to copy host and virtual machine data to a remote file server share. For a list of VMM library servers, run the following cmdlet from the Windows PowerShell command line:

```
$libraryservers = get-scvmmserver <VMM management server name> | get-sclibraryserver
```

For more information, see [Get-SCVMMServer](#) and [Get-SCLibraryServer](#).

Back up all files on library shares to a shared folder on a remote file server, including files with the following extensions:

- .vhd and .vhdx
- .iso
- .vmx
- .ps1
- .vmc
- .vsv

Back up VMM private clouds

To orchestrate and automate the replication and failover of virtual machines located in VMM clouds, you can use [Azure Site Recovery Manager](#). You can replicate in the following ways:

- From one on-premises VMM site to another, using Hyper-V replication or SAN replication.
- From an on-premises VMM site to Azure, using Hyper-V replication.

Back up registry keys, encryption keys, and credentials

Use the following guidelines to back up registry keys, encryption keys, and non-VMM managed credentials:

- **Registry keys:** VMM uses multiple registry keys to store important settings. Settings are stored in the following registry key and its subkeys: **HKLM\Software\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings**.

You should back up this entire section of the registry, with the possible exception of the SQL subkey. If you back up the SQL subkey, you capture the database name, location, and other details at the time of backup, which might not match the VMM database details that you want at the time of recovery.

To back up registry keys, you can use the Regedit **Export** function, or any other tool that is used in your environment to back up registry keys.

- **Encryption keys in Active Directory Domain Services:** If distributed key management (DKM) is configured, then you are storing VMM-related encryption keys in Active Directory Domain Services (AD DS). To back up these keys, back up Active Directory on a regular basis.
- **Non-VMM managed credentials:** Some credentials that are related to VMM are managed by the Windows Credential Manager on the VMM management server. To access the Credential Manager, in

Control Panel, select **All Control Panel Items**, and then click **Credential Manager**. Click **Back up Credentials** to back up any VMM-related credentials.

Back up non-Microsoft user interface add-ins and other non-Microsoft applications

You can use non-Microsoft user interface (UI) add-ins to extend the functionality of the VMM console. The data that is used by a UI add-in might be stored on the local server or on a remote computer, and it might be configured with a specific set of permissions. Consult the backup guidelines of your specific UI add-in.

For any other non-Microsoft applications, refer to the applications' specific backup guidelines.

Restore the VMM environment

Restore the VMM database if necessary

If the VMM database must be restored, restore it first, using the process that corresponds to your backup method. For example, to restore using SQL Server tools, see [Restore and Recovery Overview \(SQL Server\)](#).

If the VMM database is the only element that you need to restore, and you want information about the **SCVMMRecover.exe** tool, see [Backup-SCVMMServer](#).

Restore the VMM server on the same computer

If you are using the same computer for the VMM server, perform a system state restore on that computer (otherwise, skip this section). If you do this, the SID of the VMM server remains the same, and fewer steps are required to restore your VMM environment.

After you have restored the VMM server, take the following steps:

1. Remove any hosts or virtual machines from the VMM console that were removed after the last backup. If a host has been removed after the last backup, then it appears as "Not Responding" and all virtual machines on the host appear as "Host Not Responding". If the host is present but a virtual machine has been removed after the last backup, then the virtual machine appears as "Missing."
2. Add any hosts or virtual machines that were added after the last backup.

Restore the VMM server on a different computer

If you plan to restore the VMM management server onto a different physical computer, first review the hardware requirements in System Requirements.

Next, re-install VMM on the selected server, and point this VMM server to the VMM database. Because this server will have a different SID than the original computer, a few steps are necessary to bring it current with your environment. These steps include reassociating hosts with the new VMM server (otherwise, they remain mapped to the original computer's machine account).

Update hosts with the new VMM management server

1. Open the VMM console.
2. Review the lists of hosts and virtual machines, as needed, to prepare for later steps in this procedure:
 - To review the list of servers, in the **Fabric** workspace, on the left, click **Servers**.
 - To review the list of virtual machines, in the **VMs and Services** workspace, on the left, click **All Hosts**.
3. Remove any hosts or virtual machines from the VMM console that were removed after the last backup. If a host has been removed after the last backup, then it appears as "Not Responding" and all virtual machines on the host appear as "Host Not Responding". If the host is present but a virtual machine has been removed

after the last backup, then the virtual machine appears as "Missing."

4. Add any hosts or virtual machines that were added after the last backup.
5. Identify managed computers that are marked as "Access Denied," right-click each one, click **Reassociate**, and then provide the administrative credentials.
6. If you are restoring a VMM management server that was also a library server, then the new computer lists the original VMM server as the default library server. From the **Library** view, remove the original library server, and then add the new computer as a library server.

You might also need to reassociate servers in the perimeter network (also known as DMZ, demilitarized zone, and screened subnet), as described in the next section.

Reassociate servers in a perimeter network

After you restore a VMM server, servers on a perimeter network might initially appear as "Not Responding." In that case, perform the following steps.

1. Sign in to each server on the perimeter network, and then locate the VMM account. The VMM account is a local administrator account with a 10-character user name of **scvmm** plus 5 random characters.
2. Change the password of the VMM account on each server.
3. On the VMM management server, in the **Host Properties** dialog box, click the **Options** tab, and then assign each server the same password that you created in step 2.

Restore VMM library servers

To restore a library server after data loss, restore the file server shares, and then restore the data back onto the shares.

After you restore the VMM management server and the VMM database, library servers are listed in the VMM console. As needed, reassociate these listings with the physical library servers.

1. If the newly restored computer has the same name as the original computer, install the Virtual Machine Manager agent locally on that computer and then reassociate that computer with the VMM management server.
2. If the newly restored computer has a different name than the original computer, use the VMM console to remove the original computer from the list of managed computers, and then add the new computer.

Restore registry keys, Active Directory objects, and non-VMM managed credentials

Use the following guidelines to restore registry keys, Active Directory objects, and non-VMM managed credentials:

- **Registry keys:** To restore registry keys that were previously backed up, you can use the Regedit **Import** function or any other tool that is used in your environment to back up and restore registry keys. However, do not restore the SQL subkey if the database name, location, and other details that it contains do not match what you want for the VMM database at the time you are restoring the registry keys.
- **Active Directory objects:** If distributed key management (DKM) is enabled in your VMM environment, VMM stores some data in Active Directory, such as RAA passwords, product key information, and Virtual Machine Role data. After re-installing VMM, if needed, you can re-enter some of the data that was stored in Active Directory, such as RAA passwords and product key information. After you re-install VMM and (if necessary) restore Active Directory, the data in Active Directory continues to be accessible to VMM.
- **Non-VMM managed credentials:** In Control Panel, select **All Control Panel Items**, and then click **Credential Manager**. Click **Restore Credentials** to restore any VMM-related credentials that were

previously backed up.

Post-restore tasks

Depending on your VMM configuration, you might need to do some of the following tasks after you restore your VMM environment:

Configure Always On Availability Groups

If the VMM database was configured by using SQL Server Always On Availability Groups, you must complete a few tasks to ensure that the database is correctly configured with an availability group.

Reinstall Windows Azure Pack

If Windows Azure Pack (WAP) was deployed in your environment to support tenants by using VMM, then you'll have to reinstall it after you restore the VMM environment. For more information about Windows Azure Pack for Windows Server, see [Windows Azure Pack for Windows Server](#).

Install additional VMM consoles

If you had to replace any servers on which VMM consoles were installed, re-install the consoles on those servers.

Update virtual machine templates

All virtual machine templates that were restored must correctly specify the virtual hard disk that contains the operating system.

1. In the VMM console, open the **Library** workspace, expand **Templates**, and then click **VM Templates**.
2. In the **Templates** pane, right-click the virtual machine template that you want to update, click **Properties**, and click the **Hardware Configuration** page to update the settings.

Restore Windows Azure Hyper-V Recovery Manager

If Windows Azure Hyper-V Recovery Manager is implemented in the VMM environment, then you must perform a few steps to restore the Windows Azure Hyper-V Recovery Manager Provider.

Review add-ins, driver packages, and certificates

After you restore VMM, review the following items to ensure that you have taken necessary steps for your add-ins, driver packages, and certificates:

- **Non-Microsoft user interface add-ins:** To restore any non-Microsoft user interface add-ins or any other non-Microsoft party applications, consult the respective application's restore guidelines.
- **Driver packages:** Driver packages that were previously added to the VMM library might not be discovered correctly after a restore. They might have to be removed and be re-added.
- **Certificates:** Any VMM-related certificates on hosts must be updated with the information of the new VMM management server.

NOTE

After you re-install VMM, VMM updates the account control lists (ACLs) that became outdated due to the failure. No further intervention is required.

Set up the library in the VMM compute fabric

12/7/2018 • 3 minutes to read

Read this article to understand the System Center - Virtual Machine Manager (VMM) library and how you interact with it.

The VMM library is a file share that includes a catalog of resources that are used to deploy virtual machines and services in the VMM fabric. The library stores:

- **File-based resources** such as virtual hard disks, ISO images, and scripts, driver files and application packages (SQL Server data-tier applications and Web Deploy).
- **Non-file-based resources** such as virtual machine templates and service templates that are used to create VMs and services.
- **Offline virtual machines** are stored in the library.

When you install VMM by default a single library share is created on the VMM management server. You can add additional shares. For high availability you can deploy a failover cluster of file servers. Scale out file server (SOFS) isn't supported. You interact with libraries and library resources using the **Library** view in the VMM console.

What can I do in the library?

RESOURCE TYPE	WHAT CAN I ADD?
File-based resources	Virtual hard disks (.vhd/.vhdx/.vmdk), ISO image files (.iso) PowerShell scripts (.ps1), SQL Server scripts (.sql), Web Deploy (MSDeploy) packages (.zip), SQL Server data-tier apps - DACS (.dacpac), Server App-V packages (.osd), driver files (.inf), answer files (.inf, xml), virtual floppy disks (.vfd/.flp)

RESOURCE TYPE	WHAT CAN I ADD?
Templates and profiles	<p>Templates help you to quickly create VMs and services with consistent settings. You create and add profiles with specific settings to templates.</p> <p>VM templates are used to create a single VM. A template can be created from an existing virtual hard disk, another VM template in the library, or from a VM deployed on a host.</p> <p>Service templates are used to create multiple VMs and can include settings for Windows Server roles and features. In addition to using hardware and guest OS profiles service templates can leverage application and SQL Server profiles.</p> <p>Hardware profiles define hardware settings such as CPU, memory, priority of VM for resource allocation on host.</p> <p>Guest OS profiles define operating system settings that will be applied when a VM is created from a template.</p> <p>Application profiles provide instructions required to install an app. VMM supports these mechanisms for app deployment - data-tier apps (DAC) and WebDeploy (MSDeploy), running a script created for Windows installer (.msi), setup.exe, Windows PowerShell Desired State Configuration (DSC), Puppet, and Chef.</p> <p>SQL Server profiles provide instructions for customizing a SQL Server instance for a SQL Server DAC.</p>
Equivalent objects	<p>Equivalent objects are user-defined groupings of library resources that are considered equivalent. After you've marked objects as equivalent, when you point to a specific virtual disk on a library share in a template or profile, VMM can substitute any equivalent object when a VM or service is created. This means you can author templates and profiles without relying on specific physical resources, and resources can be serviced without affecting the availability of templates and profiles. VMM supports virtual disks, .iso images, and custom resources as equivalent objects.</p>
Cloud libraries	<p>Read-only library shares that are assigned to a private cloud and a stored node where self-service users with appropriate permissions can store VMs and services. You can add resources to cloud libraries to make them available for private cloud users.</p>
Self-service user content	<p>Self-service users can upload their own resources which can be used when they author templates. Users can share resources with other self-service users.</p>
Stored VMs and services	<p>Users can store their VMs that aren't in use in the stored node of the cloud library.</p>
Update catalog and baselines	<p>If you manage updates through VMM then WSUS update baselines are stored in the library.</p>

RESOURCE TYPE	WHAT CAN I ADD?
Custom resources	Add custom resources so that resources that would otherwise not be indexed show up in the library. To do this you create a folder with a .CR extension and save it to a library share. Folder contents are available to all users who can access the share. Examples for custom resources includes pre and post-execution scripts, and custom installation packages.

Next steps

[Learn about](#) adding file-based resources to the library.

Add file-based resources to the VMM library

3/14/2019 • 2 minutes to read

After you've [set up](#) the System Center - Virtual Machine Manager (VMM) library, use this article if you want to add file-based resources to the library, and mark objects in the library as equivalent.

You can add file-based resources to the library as follows:

- Copy files to the share in the VMM console
- Import and export file-based resources between library shares
- Copy files to the library share from outside the VMM console

NOTE

Sysprep the virtual hard disk and then add it to the VMM library [Learn more](#).

Copy files to the share in the VMM console

1. Go to **Library > Library Servers**.
2. Right-click a library share > **Explore**.
3. Copy files to the share.

Import and export files between libraries

1. Click **Library > Import Physical Resource**.
2. Select whether to import a resource or custom resources, select the destination library server, share, and optionally a folder. Click **OK > Import**. Verify the import in **Library Servers > target location > Physical Library Objects**.
3. To export click **Export Physical Resource**.
4. Right-click a library share > **Explore**. Select the resources you want to export (Select and hold the SHIFT key for multiple) and click **OK**. Select a destination folder and click **OK > Export**.
5. Copy files to the share.

Mark objects as equivalent

You can group library resources together so they're considered equivalent. Then when you create templates and profiles and point to a specific virtual disk on a library share, VMM can substitute any equivalent object when a VM or service is created. This means you can author templates and profiles without relying on specific physical resources, and resources can be serviced without affecting template and profile availability.

VMM supports virtual disks, .iso images, and custom resources as equivalent objects. Equivalent resources must be the same file type.

You'll need to be an admin, delegated admin, or self-service user to mark objects as equivalent. Delegated admins can mark on library shares within their scope. Self-service users can mark objects in their user role data path.

1. Click **Library > Library servers**.
2. For admins and delegated admins the **Library Server** column indicates the location of each resource. Self-services users should expand **Self Service User Content > Type** to sort library resources.

3. Right-click the resources > **Mark Equivalent**.
4. In **Equivalent Library Objects**, type the family name and release value to create a new equivalent set, or click a family name to add to an existing set. Objects must have the same family name, release value, and namespace (automatically assigned by VMM) to be equivalent.

Next steps

[Learn about](#) adding profiles to the VMM library.

Add profiles to the VMM library

12/7/2018 • 10 minutes to read

Use this article to learn about System Center - Virtual Machine Manager (VMM) profiles, and how to add them to the VMM library.

A VMM profile contains settings that are used when you create a new virtual machine or virtual machine template. Profiles make deployment easier by helping you to quickly create VMs with consistent settings. Profiles can be used to restrict the settings that are available to self-service users who create new VMs.

PROFILE	DETAILS	USED FOR VM TEMPLATES	USED FOR SERVICE TEMPLATES
Hardware profile	Defines hardware configuration settings, such as CPU, memory, network adapters, a video adapter, a DVD drive, and the VM priority when resources are allocated on a VM host.	Yes	No
Guest operating system profile	Defines operating system configuration settings that are applied to a VM, including the type of operating system, the computer name, administrator password, domain name, product key, time zone, answer file, and RunOnce file.	Yes	No
Application profile	Provides instructions for installing an application. VMM supports multiple mechanisms for application deployment. Two of these mechanisms are for specific application packaging technologies: data-tier applications (DAC) and WebDeploy (MSDeploy). A third mechanism enables you to install any application by running a script. You can use scripts that are created for Windows Installer (MSI), Setup.exe installation programs, Windows PowerShell Desired State Configuration (DSC), Puppet software, and Chef software.	No	Yes

PROFILE	DETAILS	USED FOR VM TEMPLATES	USED FOR SERVICE TEMPLATES
SQL Server profile	Provides instructions for customizing an instance of Microsoft SQL Server for a SQL Server DAC when a virtual machine is deployed as part of a service.	No	Yes
Capability profile	Defines limits and capabilities for a specific set of resources, for example settings for network adapters, processor ranges, and memory. Capability profiles are used for hardware profiles or in cloud deployment. For example, you could configure a private cloud and assign it a Hyper-V capability profile that requires all resources to be highly available. In this example, you need to set up library resources such as hardware profiles to align with the capability. To learn more, review this article .	Yes	Yes
Physical computer profile	Defines settings used to provision servers	No	No

Create a hardware profile

1. In the VMM console click > **Library** > **Create** > **Hardware Profiles**.
2. In **New Hardware Profile** > **General**, type in a profile name. You can create a hardware profile with the default settings but you probably want to tailor them. In **Hardware Profile** you can specify the hardware settings.
3. In **Compatibility** you can specify that a capability profile should be assigned to the hardware profile. Remember, capability profiles help limit available options when you're creating a new VM.
4. In **General** you can define how many virtual processors to assign to the VM. You can specify memory, startup and dynamic memory range. Startup specifies the memory that is assigned to the VM during startup. After startup this memory can be reclaimed back from the VM in accordance with the minimum memory settings.
5. In **Bus Configuration** you add and remove hardware that supports storage device.
6. In **Network Adapters**, you specify the number of network adapters in the VM, whether they'll have a static IP address or an address allocated from a pool, the MAC address and port profile. The port profile can be used to control how bandwidth is used on the adapters.
7. In **Advanced** you can specify high availability and performance settings. In **Availability** specify whether the VM should be highly available in deployed in a cluster. In **BIOS** select the order of virtual device and when **Num Lock** is enabled for password entry. In **CPU Priority** specify relative priority of CPU usage for the VM. If you set to High, the VM has more access to resources than those set to Low. In **Virtual NUMA** specify when the VM can use [virtual NUMA](#). In **Memory Weight** specify the relative memory priority for the VM.
8. After you've finished creating the hardware profile, you can right-click it to configure additional properties. In **Dependencies** you see any dependencies for the profile. For example if a library-based file is required we'd see it here. In **Access** you can see the roles or users who have permissions to use this profile. In **Validation Errors**

you can check for errors.

9. After you've created the hardware profile, you can use it when you configure a virtual machine template, or create a virtual machine. You can select a complete hardware profile, or select it and then tweak settings for the individual VM or template.

Create a guest OS profile

1. In the VMM console click > **Library** > **Create** > **Guest OS Profiles**.
2. In **New Guest OS Profile** > **General**, type in a profile name. In **Guest OS Profile** you specify the OS settings.
3. In **General Settings** > **Operating System**, specify the VM operating system. In **Identity Information**, specify the actual machine name of the VM. You probably want a unique name, so that you can specify a wildcard to generate a new name for each VM. You can also use characters ### to set an increasing numeric value. For example, if you type in ContosoVM-## it generates machines named ContosoVM-01, ContosoVM-02 etc. In **Admin Password** specify local admin permissions require a password. you can use the predefined Run As account. In **Product Key** type in the key for the OS installation. If you add an answer file under **Scripts**, you can select the **Product key provided by answer file** settings. In **Time Zone** specify the time location for the VM.
4. In **Roles and Features** specify what should be installed on the VM. Note that this setting is only used for the profile is used in a VM template which is then used in a service template.
5. In **Networking** specify domain settings for the VM and credentials to use for joining the domain.
6. In **Scripts**, specify any scripts you want to use for the VM. Scripts must be located on the library share. For example an installation answer file. The **GUIRunOnce** option allows you to run a script the first time a user logs on to the VM.
7. After you've created the guest OS profile you can right-click it to configure additional properties. In **Dependencies** you see any dependencies for the profile. For example Run As accounts. In **Access** you see the roles or users who have permissions to use this profile.
8. After you've created the hardware profile, you can use it when you configure a virtual machine template or create a VM.

Create an application profile

1. In the VMM console click > **Library** > **Create** > **Application Profiles**.
2. In **New Application Profile** > **General**, type in a profile name. In **Application Configuration** you can specify the app settings.
3. In **Application Configuration** > **OS Compatibility**, specify the guest operating systems that are compatible with the application profile.
4. Click **Add** and select the type of application or script that you want to apply to the profile. To deploy any app type select **General**. To deploy SQL Server DAC packages or script, select **SQL Server Application Host** so that you can add packages and scripts to the profile. To deploy Web applications, click **Web Application Host** so that you can add Web Deploy packages and scripts to the profile.
5. If you selected **General** you can add more than one more of application or script to the profile.
6. For applications you can specify settings such as certificate, ports, and folders. You can also specify that deployment of the app should be managed by a script. You can specify the script name, and specify when it should run.
7. Select **Scripts** to add an unlimited number of scripts and properties such as parameters and security settings. For example you can configure script to create a guest cluster out of multiple VMs deployed by VMM. For example, you can specify that one script to run at Creation: First VM (to form the cluster on the first virtual machine) and a different script to run at Creation: VMs After First (to add additional virtual machines to the cluster).
8. After you're done can verify that the profile was created in **Library** > **Profiles** > **Application Profiles**.

9. You use application profiles in service templates. For example you could create a number of VM templates with hardware and OS profiles. Then you create a service template that includes those VM templates, and the application profiles, to create a set of VMs that are configured and deployed together as a single entity.

Create a SQL Server profile

1. In the VMM console click > **Library > Create > SQL Server Profiles**.
2. In **New SQL Server Profile > General** type in a profile name. In **SQL Server Configuration** you can specify the app settings.
3. In **Application Configuration > Add > SQL Server Deployment**. A SQL Server deployment corresponds to a single instances of SQL Server. If you want multiple instance of SQL Server on the same VM you need to create multiple deployments.
4. In **SQL Server Deployment** click **Deployment 1** and specify the deployment name, the SQL Server instance details. The RunAs account is optional, the VMM service account is used if you don't specify it.
5. In **Configuration** type in the path to the SQL Server installation file (setup.exe) and the SQL Server admins.
6. In **Service Account** specify which accounts to use.

Create a capability profile

The exact settings for a capability profile depends on the profile in use. As an example let's configure the Hyper-V capability profile to specify high availability for resources used in a private VMM cloud.

1. In the VMM console, click > **Library > Create > Capability Profiles**.
2. In **Create Capability Profile > General**, type in a profile name. In **Capabilities** you specify the profile settings.
3. In **Capabilities > Fabric Compatibility** select **Hyper-V virtualization host**. You could also elect to set up a custom capability profile.
4. Set up the hardware configuration settings for the profile. The settings are similar to those used in a [hardware profile](#). However, in capability profiles these settings represent limits rather than exact values.
5. In **Advanced > Availability**, select **Highly available VM mode > Use default > Required**.
6. Complete the wizard. After you've created the profile you can select and enable it in **VMs and Services > cloud name > Properties > Capability Profiles**.
7. Remember that all of the other profiles and templates used for VMs in the cloud need to match the capability profile requirements and the high availability setting.

Create a physical computer profile

VMM can are used to provision physical computers into Hyper-V hosts or into a scale-out file server (SOFS). When you're provisioning physical computers you can use a physical computer profile to specify settings for the machine. Create a physical computer profile as follows:

1. In the VMM console click > **Library > Create > Physical Computer Profile**.
2. In **New Physical Computer Profile > Profile Definition** type in a profile name and description.
3. In **OS Image** select a virtual hard disk from the library share. It should be running Windows Server 2012 R2 or later. To create the hard disk you can create a VM, install the guest operating system, and then use Sysprep with **/generalize** and **/oobe**. If the disk is dynamic VMM converts it to a fixed disk during deployment. We recommend that you use a fixed disk type to help protect user data and increase performance.
4. In **Hardware Configuration** set up network adapters, disks and partitions, and any drivers.
5. In **Network Adapters** click **Connectivity Properties** to set up consistent device naming (CDN) for the adapter. Specify whether to allocate an IP address with DHCP or from a static pool. If it's a physical network adapter connected to a logical switch this option isn't available.

6. In **Disk** specify the partitioning scheme for the first disk. Select Master Boot Record (MBR) for BIO. or GUID Partition Table (GPT) for EFI. Specify a volume label, what free disk space to use, and what to designate as the boot partition. VMM copies the .vhd or .vhdx file to the boot partition and automatically creates a system partition on the same disk.
7. In Driver filter specify the driver files to be applied to the operating system during deployment. You can filter drives with plug and play IDs or with specific tags. With the tag option you need to add driver files to the library and assign corresponding tags to the library share before deployment.
8. In **OS Configuration** set up the domain, the password for the local admin, name and organization, product team, time zone, and an answer file for additional setup options. In GUIRunOnce you can specify commands or scripts that should run the first time a user signs into the machine.
9. Verify the settings in **Summary** and click **Finish**. You can check that the physical computer profile in **Library > Profiles > Physical Computer Profiles**.

Next steps

Learn about creating [VM templates](#) and [service templates](#) in the VMM library, and adding profiles to them.

Add VM templates to the VMM library

3/19/2019 • 4 minutes to read

Read this article to learn about VM templates and how to manage them in the System Center - Virtual Machine Manager (VMM) library.

Templates help you to create VMs with consistent settings. VMM provides two types of templates:

- VM templates are database objects that are stored in the VMM library. They're used to quickly set up VMs.
- [Service templates](#) define how a service is configured. They include information about the VMs that are deployed as part of the service, which applications to install on VMs, and the network settings that should be used. Service templates usually contain VM templates.

There are two methods for creating a VM template:

- From an existing virtual hard disk or VM template in the VMM library.

NOTE

Ensure the virtual hard disk was Sysprepped.

- From an existing VM deployed on a host.

Before you start

- You can base a new VM template on an existing VM template or on a virtual hard disk that is stored in the library. You can configure hardware settings, guest operating system settings, application installation, and Microsoft SQL Server instances. You can configure each of these settings manually or you can import the settings from an existing profile.
- Static IP address settings are available only when you deploy a VM from a VM template.
- If you create a VM template based on Linux, some of the Linux-specific settings, such as operating system specialization, work only if you deploy the Linux-based VM on a Hyper-V host.
- The option to create a VM template that's based on an existing VM on a host isn't applicable for Linux-based virtual machine templates.
- Application deployment, SQL Server deployment, and configurable service settings only apply when you deploy a VM as part of a service.
- If you grant rights for a particular template to a user that does not have rights to the Run As account that's specified in the template, then the user can potentially extract the credentials for the Run As account from the template during deployment.
- Before creating a template based on a VM, you'll need to create a new local administrator account on that VM. Using the default built-in administrator account will cause Sysprep to fail. In addition, make sure that the VM isn't joined to a domain.

Create a VM template based on an existing VHD or VM template in the library

1. Click **Library > Create > Create VM Template**.
2. In the Create VM Template Wizard, select **Source > Use an existing VM template or virtual hard disk stored in the library**. Select the disk or template in **Select VM Template Source**.

3. In **Identity**, type in a template name and description.
4. In **Configure Hardware**, specify the hardware settings. You can select to use an existing hardware profile. Note that the profile and hardware options will depend on which you're configuring - generation 1 or generation 2 VMs. To learn more, review [how to create a hardware profile](#).
5. In **Configure Operating System**, specify the machine settings. You can use a guest OS profile or configure specific settings. To learn more, review [how to create a guest OS profile](#).
6. In **Configure Applications**, set up app settings. This isn't relevant if you're using the VM template to deploy VMs that aren't part of a service. To learn more, review [how to create an application profile](#). If you're configuring SQL Server set in up in **Configure SQL Server**, review [how to create a SQL Server profile](#).
7. In **Summary**, review the settings and click **View Script** if you want to see the script that will be used to create the template. Then click **Create**. In **Jobs**, you can track the template being created. Wait for the **Completed** status.
8. When you create a VM you'll be able to create it based on the template you've just created.

Create a VM template based on a VM deployed on a host

1. Click **Library > Create > Create VM Template**.
2. In the Create VM Template Wizard > select **Source > From an existing virtual machine that is deployed on a host**. Select the VM in **Select VM Template Source**.
3. In **Identity**, type in a template name and description. Note that the template will destroy the source VM and data on it could be lost. Clone it if it's important to you.
4. In **Configure Hardware**, click **Next**.
5. In **Configure Operating System**, specify the guest operating system settings. You can use a guest OS profile or configure specific settings. To learn more, review [how to create a guest OS profile](#).
6. In **Select Library Server**, click the library server for the VM you're using for the template and in **Select Path** specify the share/folder.
7. In **Summary**, review the settings and click **Create**. In **Jobs**, you can track the template being created. Wait for the **Completed** status.

You can create VMs based on the template you've just created.

Assign a storage QoS policy template

1. After step 3 in the [previous procedure](#), in **Configure Hardware**, click **Advanced** under **Bus configuration**, and select appropriate option under **Storage QoS Policy**.
2. Proceed with rest of the steps to complete the wizard.

Next steps

Review the following on how to [create VMs based on the template](#) you've just created.

Add service templates to the VMM library

5/30/2019 • 3 minutes to read

Read this article to learn about setting up service templates in the System Center - Virtual Machine Manager (VMM) library.

Service templates group VMs together to provide an app. They contain information about a service, including the VMs that are deployed as part of the service, the applications installed on VMs, and the network settings that should be used. You can add VM templates, network settings, applications, and storage to a service template.

Service templates can be single or multi-tier:

- A single tier service contains one VM used as a specific app.
- A multi-tier service contains multiple VMs. For example you could create a three-tier service with a backend tier running a SQL Server database, a middle tier running the business server, and a third tier running a frontend web interface.
- Tiers can be added based on a copy of an existing VM template (which can be customized) or a virtual hard disk in the library.

You set up service templates using the VMM service template designer.

Before you start

You can create service templates if you have VMM admin or delegated admin permissions, or if you have a self-service user account with **Author** enabled.

Create a service template

1. Click **Library > Create > Create Service Template**.
2. In **New Service Template > Name**, specify a template name. In **Release** indicate the template version.
3. To configure a tier using the predefined templates click in the designer workload and select a predefined tier pattern (blank, 1, 2, or 3 tiers). Click **Save and Validate** to save the template. After its created you can click on a template object to modify its name, release version or users/roles that can access it.
4. When the tier appears in the workspace drag a VM template to it. The properties of the VM template are applied to the tier. Note that this doesn't establish a link between the tier and the template. Changing the template properties does not modify the tier properties.

NOTE

You can also click **Add Machine Tier** to add a tier manually. This opens the **Create Machine Tier Template** wizard. In **Select Source** select a source for the tier. You can use an exact copy of an existing VM template, or to customize an existing VM template. Click **Browse** to select the template or hard disk. In **Additional Properties**, you'll configure the tier properties described in the next step.

5. You can click on a tier to access its properties in the details pane of the designer. Click **View All Properties** to modify all the properties in a single view. Here's what you can modify when you select to view all:
 - In **General** specify:
 - The order in which tiers are deployed and serviced. For example if you need the database tier to be running in order to run a frontend web app, you'd set the database tier to 1.
 - Whether you want to be able to add additional VMs to the tier in order to scale out (you can scale

out to 5 VM instances in a tier).

- More than one upgrade domain to minimize service interruptions when a tier is updated. VMM will update VMs in the tier according to their upgrade domains. VMM upgrade an upgrade domain at a time. It shuts down VMs in the domain, updates them, brings them online, and moves to the next domain to reduce impact.
- The creation of an availability set for the tier. The availability set helps VMs in the service remain available during maintenance. VMM tries to separate VMs in the same availability set by placing them on separate hosts.
- In **Configure Hardware** you'll see the hardware settings for the associated VM template. You can select an alternative hardware profile, or configure hardware settings manually. To learn more, review [how to create a hardware profile](#).
- In **Configure Operating System** you'll see the operating system settings for the associated VM template. you can select an alternative guest OS profile, or configure settings manually. To learn more, review [how to create a guest OS profile](#).
- In **Application Configuration** or **SQL Server Configuration** you can select an application/SQL Server profile or configure settings for a new profile. Learn more about [application](#) and [SQL Server](#) profiles.

Add a VM network to the service template

You'll need to configure network settings for a tier by connecting the tier adapters to one or more [VM networks](#). To do this you'll add a logical network component and then use the connector tool to connect it to the adapter.

1. In the Service Template Designer, click **Service Template Components > Add VM Network**.
2. When the network appears as a component, use the Connector to connect to the appropriate NIC.

Next steps

[Set up load balancing for a service tier](#)

Manage the VMM library

12/7/2018 • 4 minutes to read

Read this article to learn how to manage the System Center - Virtual Machine Manager (VMM) library by refreshing it, moving files around, and removing orphaned resources.

Refresh the library

By default VMM refreshes library shares once every hour.

- You can change the default refresh settings in **General > Library Settings > Settings > Modify**, from to 366 hours (14 days). You can also disable automatic library refreshes.
- To manually refresh, select **Library > library server or share > Refresh share**.

During a library refresh, the following occurs:

- VMM adds these file types to the **Library** view: virtual hard disks (except for those attached to a stored VM), virtual floppy disks, ISO images, answer files, PowerShell scripts. Snapshots imported into the library with Hyper-V and VMware VMs are displayed on the **Checkpoints** tab of the VM properties. The snapshot files aren't displayed.
- VMM indexes but doesn't display these file types:
 - Files associated with stored VMs (VM configuration file, attached virtual hard disks, saved state files, imported snapshots, checkpoints).
 - Files associated with VM templates.
 - Configuration files:
 - Hyper-V (.exp -export, .vsv -savedstate, .bin)
 - Virtual Server (.vmd, .vsv)
 - VMware (.vmtx, .vmx)

Transfer files

You have direct access to copy and move library files through Windows Explorer. Each file in a library share has a unique GUID and is periodically indexed during library refreshes. After a file is refreshed you can move it to any other location on a library share managed by VMM and refresh automatically tracks the file movement. After the move, file metadata is updated during the next library refresh.

In addition you can allow unencrypted file transfers to and from a library server.

- To transfer unencrypted file transfers the feature must be allowed on both the source and destination servers.
- To enable the option click **Library > Library Server** and navigate to the server. Click **Actions > Library Server > Properties**, select **Allow unencrypted file transfers**.

Disable and remove file-based resources

You can remove a file-based resource either temporarily or permanently from the library.

- To disable resources click **Library > Library servers >** and select the library share. Select the resource and click **Actions > Disable**. click **Enable** to re-enable.
- To remove files we recommend you use VMM rather than simply deleting the file resources. When you remove the file in the library any resources that use the file are updated automatically. To remove a file click **Library >**

Library servers > and select the library share. Select the resource and click **Actions** > **Remove**. Select **Yes** to confirm.

Remove a library server or share

There are circumstances in which you need to remove a library server or share. For example if you're no longer using the resources on a share or you temporarily want to remove those resources.

- To remove a library share click **Library** > **Library servers** > and select the library share. In Actions, click **Library Share** > **Remove**. Click **Yes** to confirm. Note that removing a share doesn't delete the files on it. They are no longer be indexed by library refreshes.
- To remove a library server click **Library** > **Library servers**. Click **Actions** > **Library Server** > **Remove**.
- Note that:
 - Specify an account with administrative permissions on the server.
 - VMM provides a list of dependent resources. If you process VMM removes any references to the removed files on dependent resources. When you remove a library server the **Library Server** role is removed from the VMM agent running on the server. If the server isn't performing any other VMM roles the agent is removed.
 - If you remove a highly available library server the cluster is removed from the **Library** view. The individual cluster nodes aren't removed, but they're not displayed in the library. To remove the nodes from VMM remove the VMM agent from each computer.

Remove orphaned resources

When you remove a library share from VMM management, and there are templates that reference resources that were located on the library share, a representation of the library resource appears in the VMM library as an orphaned resource.

To remove orphaned resources, modify the templates that reference the orphaned resources to use valid library resources in the VMM library. If you add the library share again, VMM doesn't automatically reassociate the template with the physical library resource. You must complete these steps to correct template issues and to remove any orphaned resources.

1. Click **Library** > **Orphaned Resources**.
2. You won't be able to delete an orphaned resource until templates that reference it are updated to valid references. To view the templates right-click the orphaned resource > **Properties**. To update the template click it and then in the **Properties** dialog, locate the resource that's missing > **Remove**.
3. Add a new resource that's valid.
4. When you've completed these steps for all templates, closed the **Properties** dialog. To verify there aren't any dependencies right-click the orphaned resource > **Properties** > **Dependencies**. Then right-click the orphaned resource > **Delete**.

Set up host groups in the VMM compute fabric

12/7/2018 • 2 minutes to read

Read this article to learn about setting up and managing host groups in the System Center - Virtual Machine Manager (VMM) fabric.

A VMM host group is a logical entity that groups fabric resources together. You can group virtual machine hosts or clusters, or create nested host groups. After you've created host groups you assign and configure resources at the host group level. Those resources are then applied to all hosts and clusters in the group.

You can create host groups based upon different criteria. For example based on physical location, hardware capabilities, or specific workloads. You can assign permissions for host groups to the VMM admin, delegated administrators, and read-only admin user roles. Members of these user roles can view and manage the fabric resources that are assigned to them at the host group levels. When you create private clouds in VMM, you select which host groups will be included in the cloud, and then allocate resources in the host groups to the cloud.

Create host groups

1. Click **Fabric > Servers > All Hosts > Create Host Group**.
2. Type in a group name. To create a host group at a specific location in the tree, right-click the desired parent node, and then click **Create Host Group**.

After you've created a host group you can modify the following properties for the group.

TAB	PROPERTY
General	Configure the host group name, the location in the host group hierarchy, the description, and whether to allow unencrypted BITS file transfers.
Placement Rules	VMM automatically identifies the most suitable host to which you can deploy virtual machines. However, you can specify custom placement rules. By default, a host group uses the placement settings from the parent host group.
Host Reserves	<p>Host reserve settings specify the amount of resources that VMM sets aside for the host operating system to use. For a virtual machine to be placed on a host, the host must be able to meet the virtual machine's resource requirements without using host reserves. You can set host reserves for individual host groups and for individual hosts. The host reserve settings for the root host group, All Hosts, sets the default host reserves for all hosts. You can configure reserve values for the following resources:</p> <p>CPU</p> <p>Memory</p> <p>Disk I/O</p> <p>Disk space</p> <p>Network I/O</p>

TAB	PROPERTY
Dynamic Optimization	Configure dynamic optimization and power optimization settings. Dynamic optimization balances the virtual machines load within a host cluster. Power optimization enables VMM to evacuate hosts of a balanced cluster and turn them off to save power.
Network	<p>View inheritance settings, and configure whether to inherit network logical resources from parent host groups. The network logical resources include the following:</p> <p>IP address pools</p> <p>Load balancers</p> <p>Logical networks</p> <p>MAC address pools</p>
Storage	View and allocate storage to a host group.
Custom Properties	<p>Manage custom properties for the following object types:</p> <p>Virtual machine</p> <p>Virtual machine template</p> <p>Host</p> <p>Host cluster</p> <p>Host group</p> <p>Service template</p> <p>Service instance</p> <p>Computer tier</p> <p>Cloud</p>

Next steps

After you've created host groups, you can deploy [Hyper-V](#) hosts in the VMM fabric.

Add Windows servers as Hyper-V hosts or clusters in the VMM compute fabric

12/7/2018 • 10 minutes to read

Read this article to learn about adding an existing Windows server as a Hyper-V host server or cluster to the System Center - Virtual Machine Manager (VMM) fabric, and configuring the host and cluster properties.

The article is relevant for adding Windows server computers with or without the Hyper-V role. If you add a Windows server that doesn't have Hyper-V installed, VMM will install the Hyper-V role as long as the server meets the prerequisites.

Before you start

The prerequisites for adding an existing Hyper-V host server or cluster depend on whether Hyper-V is installed and where the server is located.

HOST LOCATION	PREREQUISITE
Server without Hyper-V	<p>If you want to add a server that doesn't have Hyper-V installed it must meet the prerequisites for Hyper-V installation.</p> <p>The server must be running supported version of Windows Server.</p> <p>If you want to add the VMM management server as a managed Hyper-V host the Hyper-V role must be installed on the server before you add it. You can't add a highly available VMM server as a managed Hyper-V host cluster.</p> <p>If you want to add a Hyper-V cluster the instructions in this article presume that the cluster already exists. Read this article if you want to create a cluster from existing Hyper-V hosts in the VMM fabric.</p> <p>This article assumes that the server you want to add already has an operating system running on it. If you want to add a bare-metal computer as a Hyper-V host or cluster read this article.</p>

HOST LOCATION	PREREQUISITE
<p>Same domain as VMM server, or two-way trusted domain</p>	<p>You must specify account credentials for an account that has administrative rights on the computers that you want to add. You can enter a user name and password or specify a RunAs account.</p> <p>If you use Group Policy to configure Windows Remote Management (WinRM) settings, note these settings:</p> <p>WinRM Service settings must be configured through Group Policy, and can only apply to hosts that are in a trusted Active Directory domain. Specifically, VMM supports the configuration of the Allow automatic configuration of listeners, Turn On Compatibility HTTP Listener, and Turn on Compatibility HTTPS Listener Group Policy settings. VMM does not support other WinRM Service policy settings.</p> <p>If you enable the Allow automatic configuration of listeners policy setting, you must configure it to allow messages from any IP address. In other words, in the policy setting, the IPv4 filter and IPv6 filter (depending on whether you use IPv6) must be set to *****.</p> <p>WinRM Client settings cannot be configured through Group Policy. These policy settings might override client properties that VMM requires for the VMM agent to work correctly.</p> <p>If you enable any unsupported WinRM Group Policy settings, installation of the VMM agent may fail.</p>
<p>Untrusted domain</p>	<p>VMM doesn't support configuring Windows Remote Management (WinRM) Group Policy settings (Service or Client) on hosts that are in an untrusted Active Directory domain. If WinRM Group Policy settings are enabled, installation of the VMM agent—required on the hosts might fail.</p> <p>In an untrusted domain, when VMM installs the agent on the servers or clusters, it also generates a certificate. The certificate is used to help secure communications with the host. When VMM adds the host or cluster, the certificate is automatically imported into the VMM management server trusted certificate store.</p>

HOST LOCATION	PREREQUISITE
<p>Disjoined namespace (DNS suffix doesn't match the domain of which it's a member)</p>	<p>The System Center Virtual Machine Manager service must be running as the local system account or as a domain account that has permission to register a Service Principal Name (SPN) in Active Directory.</p> <p>When you try to add a computer that is in a disjoined namespace, VMM checks Active Directory to see if an SPN exists. If it doesn't VMM tries to create one. If the permissions are OK VMM adds the missing SPN automatically. Otherwise, host addition fails and you'll need to add the SPN manually. To do this type: setspn -A HOST/. For example, setspn -A HOST/hypervhost03.contosocorp.com hypervhost03.</p> <p>If the host cluster is in a disjoined namespace and the VMM management server is not, add the DNS suffix for the host cluster to the TCP/IP connection settings on the VMM management server.</p> <p>If you use Group Policy to configure Windows Remote Management (WinRM) settings, review the following requirements:</p> <p>WinRM Service settings must be configured through Group Policy, and can only apply to hosts that are in a trusted Active Directory domain. Specifically, VMM supports the configuration of the Allow automatic configuration of listeners, Turn On Compatibility HTTP Listener, and Turn on Compatibility HTTPS Listener Group Policy settings. VMM does not support other WinRM Service policy settings.</p> <p>If you enable the Allow automatic configuration of listeners policy setting, you must configure it to allow messages from any IP address. In other words, in the policy setting, the IPv4 filter and IPv6 filter (depending on whether you use IPv6) must be set to *.</p> <p>WinRM Client settings cannot be configured through Group Policy. These policy settings might override client properties that VMM requires for the VMM agent to work correctly.</p> <p>If you enable any unsupported WinRM Group Policy settings, installation of the VMM agent may fail.</p>
<p>Perimeter network or workgroup</p>	<p>You'll need to install the VMM agent locally on the target host. To do this run VMM setup as an administrator and click Optional Installations > Local Agent. In Security File Folder select This host in on a perimeter network and enter an encryption key. In Host network name specify how the VMM server will contact the host server and note the computer name or IP address. Finish the wizard.</p> <p>Check that a file SecurityFile.txt is located on the VMM server. By default it's located in C:\Program Files\Microsoft System Center version\Virtual Machine Manager.</p>

Add servers

1. In the VMM console open **Fabric > Servers**.
2. Click **Add group > Add Resources > Hyper-V hosts and Clusters**.

3. In the **Add Resource Wizard** > **Resource location**, select where the server you want to add is located.
 - If you're adding a host in a perimeter network select **Windows Server computer in a perimeter network**.
4. In **Credentials**, specify credentials for a domain account that has administrative permissions on all hosts that you want to add. (For computers in an untrusted domain, you must use a RunAs account.)
5. In **Discovery scope** specify:
 - **Same domain or domains with two-way trust:**
 - If you click **Specify Windows Server** computers by names, in **Computer names** enter names or IP addresses, one per line. If you are adding a Hyper-V host cluster, specify the name or IP address of the cluster or of any cluster node.
 - If you click **Specify an Active Directory** query to search for Windows Server computers, you can type or generate a query.
 - **Untrusted domain:** Discovery page doesn't appear
 - **Disjoined namespace:** Enter the host FQDN and select **Skip AD verification**.
6. In **Target resources**, specify the computers you want to add. Repeat for all hosts. If discovery succeeds the host will be listed under **Computer name**. Add as follows:
 - **Trusted domain or disjoined namespace:** Select the check box next to each computer that you want to add, and then click Next. If you specified a cluster name or cluster node in the previous step, select the check box next to the cluster name. (The cluster name is listed together with the associated cluster nodes.)
 - **Untrusted domain:** Enter the FQDN or IP address of the server or cluster that you want to add, and then click Add. For a cluster, you can enter an FQDN or IP address of the cluster or of one of the cluster nodes.
 - **Perimeter network/workgroup:** Enter the NETBIOS name or IP address of the host in the perimeter network. Enter the encryption key you created when you installed the agent on the host, and in Security file path enter the path to the SecurityFile.txt file.
7. In the **Host settings** > **Host group** list, click the host group to which you want to assign the host or host cluster. If the host is already associated with a different VMM management server, select **Reassociate this host with this VMM environment**. If the host was associated with a different VMM management server, it will stop working on that server.
 - For a standalone host, in **Add the following path**, enter a path on the host for storing files for virtual machines that are deployed on the host, and then click **Add**. Repeat to add more than one path. If the path doesn't exist it's created automatically. If you leave the box empty, the default is %SystemDrive%\ProgramData\Microsoft\Windows\Hyper-V. As a best practice don't add default paths that are on the same drive as the operating system files.
 - For a cluster don't specify default virtual machine paths. VMM automatically manages the paths that are available for virtual machines, based on the shared storage that is available to the host cluster
8. On the **Summary** page, confirm the settings, and then click **Finish**. The **Jobs** dialog box appears to show the job status. Wait for a Completed status. Verify that the host or cluster was added in the host group > host or cluster name. The status should be **OK**.

Configure properties for Hyper-V hosts

After you've added Hyper-V hosts and servers in the VMM fabric there are a number of properties you can configure for standalone hosts and clusters.

TAB	SETTINGS
General	<p>View identity and system information for the host. This includes information such as processor information, total and available memory and storage, the operating system, the type of hypervisor, and the VMM agent version.</p> <p>Enter a host description.</p> <p>Configure whether the host is available for placement.</p> <p>Configure the remote connection port. By default, the port is set to 2179.</p>
Hardware	View or modify settings for CPU, memory, graphics processing units (GPUs), storage (including whether the storage is available for placement), network adapters, DVD/CD-ROM drives and Baseboard Management Controller (BMC) settings.
Status	<p>Lists health status information for the host. Includes areas such as overall health, Hyper-V role health and VMM agent health. In the Status pane, you can also do the following:</p> <p>View error details.</p> <p>Refresh the health status.</p> <p>Click Repair all. VMM will try to automatically fix any errors.</p>
Virtual Machine Paths/Virtual Machines	Shows the virtual machines that reside on the host, together with status information. Also enables you to register virtual machines on the host.
Reserves	Enables you to override host reserve settings from the parent host group, and configure reserved resources for the host. Configurable resources include CPU, memory, disk space, disk I/O and network capacity.
Storage	Shows storage allocated to a host, and enables you to add and remove storage logical units or file shares.
Virtual Switches	Enables you to configure virtual switches.
Placement Paths/Placement	Enables you to configure the default virtual machine paths and default parent disk paths that will be used during virtual machine placement on the host.
Servicing Windows	Enables you to select servicing windows.
Custom Properties	Enables you to assign and manage custom properties.

Properties for Hyper-V clusters

TAB	SETTINGS
-----	----------

TAB	SETTINGS
General	<p>View the name, host group and description. You can also configure the Cluster reserve (nodes) setting, and view the cluster reserve state.</p> <p>The Cluster reserve (nodes) setting specifies the number of node failures a cluster must be able to sustain while still supporting all virtual machines deployed on the host cluster. If the cluster cannot withstand the specified number of node failures and still keep all of the virtual machines running, the cluster is placed in an over-committed state. When over-committed, the clustered hosts receive a zero rating during virtual machine placement. An administrator can override the rating and place a highly-available virtual machine on an over-committed cluster during a manual placement.</p>
Status	<p>View detailed status information for the host cluster:</p> <p>Cluster validation test runs and successes. Includes a link to the latest validation report (if available). Note that accessing the report requires administrative permissions on the cluster node where the report is located. For host clusters, you can perform an on-demand cluster validation through VMM. To do this, in the Fabric workspace, locate and click the host cluster. Then, on the Host Cluster tab, click Validate Cluster. Cluster validation begins immediately.</p> <p>Online elements in the cluster: cluster core resources, disk witness in quorum, and the cluster service on each node.</p>
Available Storage	<p>Shows available storage, that is, storage logical units that are assigned to the host cluster but are not Cluster Shared Volumes (CSV).</p> <p>You can also do the following:</p> <p>Add and remove storage logical units that are managed by VMM.</p> <p>Convert available storage to shared storage (CSV).</p>
Shared Volumes	<p>Shows the shared volumes (CSVs) that are allocated to the host cluster. You can also do the following:</p> <p>Add and remove CSVs that are managed by VMM.</p> <p>Convert CSVs to available (non-CSV) storage.</p>
Custom Properties	<p>Custom properties that you manage.</p>

Deploy and manage Nano server-based Hyper-V hosts or VMs in VMM

8/29/2019 • 5 minutes to read

You can use System Center - Virtual Machine Manager (VMM) to manage hosts and virtual machines running Nano server.

Using VMM, you can add and manage existing hosts running Nano, configure bare metal machines as Nano Server-based hosts, deploy compute clusters, and storage clusters (disaggregated and hyper-converged). You can manage Nano-based VMs, including shielded VMs.

Before you start

- For VM deployment, you need to create the Nano Server virtual hard disk outside VMM.
- You can't create a VM template from a Nano Server VM in VMM. You can create a VM template from scratch using a Nano Server virtual hard disk.
- There are some known issues when joining a Nano Server VM to a domain. If you try to join the VM to a domain with customization details in a VM template, the domain information is ignored by VMM. The VM is deployed, but doesn't join the domain. As a workaround, deploy the VM and then join it to a domain. [Learn more](#). Note that joining a physical machine to a domain during bare metal deployment works fine.

Prepare a Nano server virtual hard disk

To get started with the deployment of a Nano Server-based host or virtual machines in VMM you need to create a Nano server VHD from the Windows Server VHD. The VHD should include the VMM packages:

- Add the VMM package, **Microsoft-NanoServer-SCVMM-Package**, to ensure that the VMM agent is part of the VHD.
- Add the VMM compute package, **Microsoft-NanoServer-SCVMM-Compute-Package**, to ensure that the VHD has the Hyper-V role, and that you can manage the physical server using VMM. If you install this package, don't use the **-Compute** option for the Hyper-V role).
- For the File Server role, use **Microsoft-NanoServer-Storage-Package**, along with **Microsoft-NanoServer-SCVMM-Package**.
- For a hyperconverged deployment, use **Microsoft-NanoServer-Storage-Package**, along with **Microsoft-NanoServer-SCVMM-Package** and **Microsoft-NanoServer-SCVMM-Compute-Package**.

Create a virtual hard disk for a physical machine

1. Copy **NanoServerImageGenerator.psm1** and **Convert-WindowsImage.ps1** from the \NanoServer folder in the Windows Server ISO, to a folder on your hard drive.
2. Start Windows PowerShell as an administrator. Navigate to the folder in which you placed the scripts.
3. Import the **NanoServerImageGenerator** script by running:

```
Import-Module NanoServerImageGenerator.psm1 -Verbose
```

4. Create a VHD that includes the VMM packages. To do this, run the following command which will prompt you for an administrator password for the new VHD:


```
New-NanoServerImage -MediaPath <path to root of media> -BasePath .\Base -TargetPath  
.\NanoServerVM\NanoServerVM.vhdx -ComputerName <computername> -OEMDrivers -Package Microsoft-NanoServer-  
SCVMM-Package,Microsoft-NanoServer-SCVMM-Compute-Package
```

For example:

```
New-NanoServerImage -MediaPath F:\ -BasePath .\Base -TargetPath .\Nano1\NanoServer.vhd -ComputerName  
Nano-srv1 -OEMDrivers -Clustering -EnableRemoteManagementPort -Packages Microsoft-NanoServer-SCVMM-  
Package,Microsoft-NanoServer-SCVMM-Compute-Package
```

- This example creates a VHD from an ISO mounted as F:
 - When creating the VHD, it uses a folder called Base in the same folder in which you ran **New-NanoServerImage**
 - It places the VHD in a folder called **Nano1**, in the folder from which the command runs.
 - The computer name in this example is **Nano-srv1**. It includes the OEM drivers installed for most common hardware, and has the clustering feature enabled.
 - The VHD has remote management of the Nano server enabled, even from the systems which are not in the same subnet.
 - If the server uses UEFI to boot, you need to change the script from **NanoServer.vhd** to **NanoServer.vhdx**.
5. Log in as an administrator on the physical server on which you want to run the Nano Server VHD.
 6. Copy the VHD that the script creates to the physical computer, and configure it to boot from the new VHD, as follows:
 - Mount the generated VHD.
 - Run **bcdboot d:\windows** (in this example, it's mounted under D:)
 - Unmount the VHD.
 7. Boot the physical computer into the Nano Server virtual hard disk.
 8. Log on to the Nano server Recovery Console using the administrator name and password you supplied when running the script, and obtain the IP address of the Nano server-based host. [Learn more](#).
 9. Ensure that the Nano server is joined to the same domain as the VMM server. [Learn more](#).
 10. Ensure that the VMM service account, and the Run As account, are added to the administrators group on the Nano server.

Install the VMM packages offline on an existing Nano Server

If you didn't add the VMM packages when you created the Nano Server VHD, you can install them later, as follows:

1. Copy the VHD/VHDX to a location on a Windows Server 2016 machine. For example: C:\MyNano.vhd.
2. Use PowerShell to install and import the NanoServerPackage provider of the PackageManagement (OneGet) PowerShell module:

```
Install-PackageProvider NanoServerPackage  
Import-PackageProvider NanoServerPackage
```

3. After the provider is installed, you can search and install the VMM packages (VMM agent and Hyper-V), on the VHD using these cmdlets, where **C:\MyNano.vhd** is the location of the Nano Server based VHD.

```
Find-NanoServerPackage
Install-NanoServerPackage -Name Microsoft-NanoServer-SCVMM-Package -culture en-US -ToVhd "C:\MyNano.vhd"
Install-NanoServerPackage -Name Microsoft-NanoServer-SCVMM-Compute-Package -culture en-US -ToVhd
"C:\MyNano.vhd"
```

Install the VMM packages on a running Nano server host

We recommend offline installation of the VMM packages (when creating the VHD). If you do need to install them online when the Nano server is running, do the following:

1. Copy the **Packages** folder from the local installation media to the running Nano server. For example, to C:\packages.
2. Use remote PowerShell to log onto the Nano server.
3. Add the VMM packages using the below commands:
 - To install Microsoft-NanoServer-SCVMM-Package

```
dism /online /Add-package /PackagePath:C:\packages\en-US\Microsoft-NanoServer-SCVMM-Package_en-us.cab
```

NOTE

Make sure that the en-us (Microsoft-NanoServer-SCVMM-Package_en-us.cab) and neutral (Microsoft-NanoServer-SCVMM-Package.cab) .cab files are in the same folder so that both are installed.

- To install Microsoft-NanoServer-SCVMM-Compute-Package:

```
dism /online /Add-package /PackagePath:C:\packages\en-US\Microsoft-NanoServer-SCVMM-Compute-Package_en-us.cab
```

4. Check that the VMM packages, and the associated language packs, are installed correctly by running the following command:

```
dism /online /get-packages
```

5. You should see **Package Identity : Microsoft-NanoServer-SCVMM-Feature-Package~31bf3856ad364e35~amd64~~ 10.0.14300.1003** listed twice. Once for **Release Type : Language Pack**, and once for **Release Type : Feature Pack**. The same applies for the Microsoft-NanoServer-SCVMM-Compute-Package.
6. Restart the Nano Server host.

Add the Nano server host to the VMM fabric

After the Nano server host is ready, add it to the VMM fabric. [Learn more](#).

Create a Nano server VM

To create Nano server-based VMs, you need to add few VM-specific packages to the VHD. Create the VHD for a VM as follows:

1. Copy **NanoServerImageGenerator.psm1** and **Convert-WindowsImage.ps1** from the \NanoServer

folder in VMM, to a folder on your hard drive.

2. Start Windows PowerShell as an administrator, and navigate to the script folder.
3. Import the **NanoServerImageGenerator** script with **Import-Module NanoServerImageGenerator.psm1 -Verbose**.
4. Create a VHD that includes the SCVMM packages by running the following command. You'll be prompted for an administrator password for the new VHD.

```
New-NanoServerImage -MediaPath <path to root of media> -BasePath .\Base -TargetPath  
.\NanoServerVM\NanoServerVM.vhdx -ComputerName <computername> -GuestDrivers -Package Microsoft-  
NanoServer-SCVMM-Package,Microsoft-NanoServer-SCVMM-Compute-Package
```

Example:

```
New-NanoServerImage -MediaPath F:\ -BasePath .\Base -TargetPath .\Nano1\Nano.vhd -ComputerName Nano1 -  
GuestDrivers -Package Microsoft-NanoServer-SCVMM-Package,Microsoft-NanoServer-SCVMM-Compute-Package
```

5. This example creates a VHD from an ISO mounted as F.
6. When creating the VHD it will use a folder called Base in the same directory in which you ran New-NanoServerImage
7. It will place the VHD in a folder called Nano1 in the folder in which the command runs.
8. The computer name will be Nano1 and will install virtual machine drivers running Hyper-V.
9. If you want a Generation 1 virtual machine, generate a VHD image using a .vhd extension for -TargetPath. For a Generation 2 virtual machine, generate a VHDX image with the .vhdx extension for -TargetPath.
10. In VMM, create a new virtual machine and use the virtual hard disk you created.

Next steps

[Provision a VM](#)

Run a script on a remote host by using Run Script command

5/30/2019 • 2 minutes to read


System Center Virtual Machine Manager (VMM) supports the ability to execute commands remotely, on a host by using the Run Script command feature. This feature is useful in scenarios where you want to start a service or collect information from a remote host.

You can also run the script on the host by using the custom resources that you added to the VMM library. Custom resources can consist of batch files that can execute specific commands against the server, but the use of custom resources is not required.

Example scenario 1 - Start a service

The following example scenario provides information on how you can start a service **sftlist** on a VMM host using the Run Script command feature.

1. In the VMM console, on the toolbar, click **Host** and select **Run Script Command**.



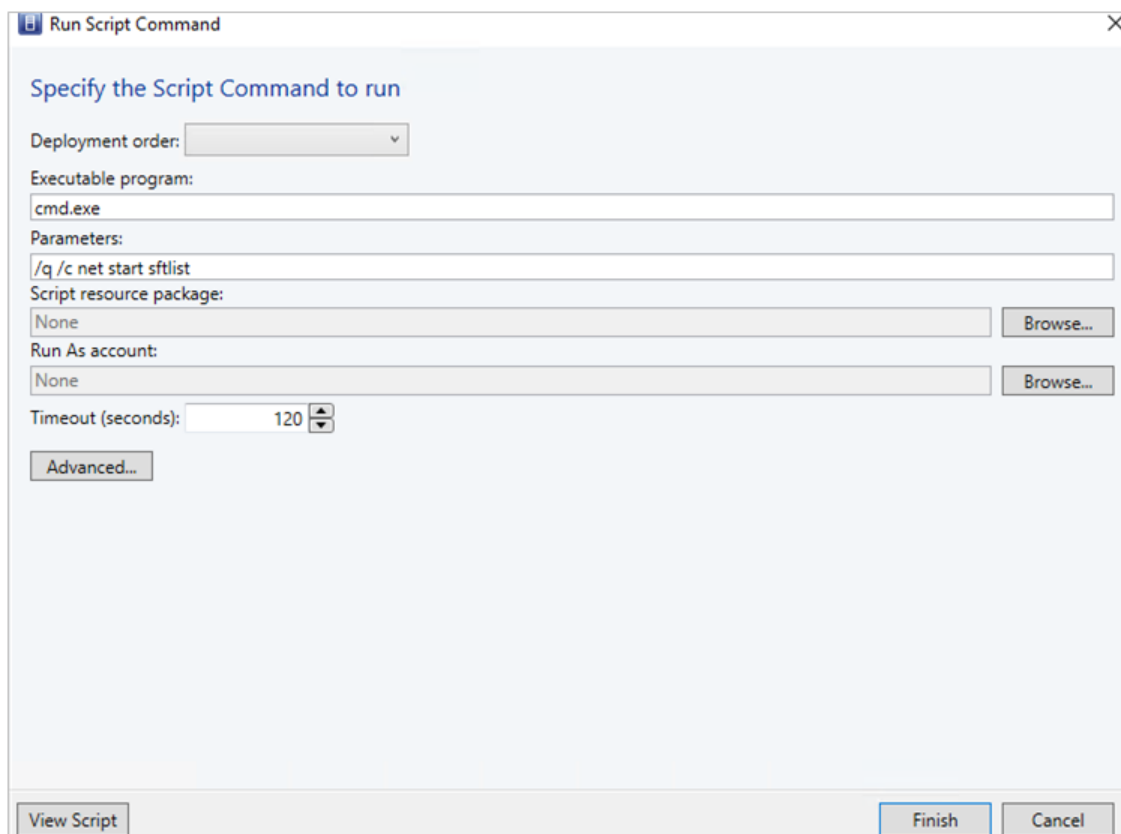
The **Run Script Command** page appears:

A screenshot of the 'Run Script Command' dialog box. The dialog has a title bar with a close button. The main area is titled 'Specify the Script Command to run'. It contains several fields: 'Deployment order:' with a dropdown arrow, 'Executable program:' with a text box, 'Parameters:' with a text box, 'Script resource package:' with a dropdown set to 'None' and a 'Browse...' button, 'Run As account:' with a dropdown set to 'None' and a 'Browse...' button, and 'Timeout (seconds):' with a text box containing '120' and a spinner. There is an 'Advanced...' button below the timeout field. At the bottom, there are three buttons: 'View Script', 'Finish', and 'Cancel'.

2. As an example, type the following details against the options displayed:

- **Deployment Order:** Select the deployment order from the drop-down menu.
- **Executable program:** cmd.exe
- **Parameters:** /q /c net start sftlist
- **Script resource package** and **Run As account:** None

The **/q** turns off the echo for the command line and **/c** carries out the command and then terminates. If you do not use these switches, the outcome will be a return code 0 from the cmd.exe process and will not display the actual result of the command in the job.



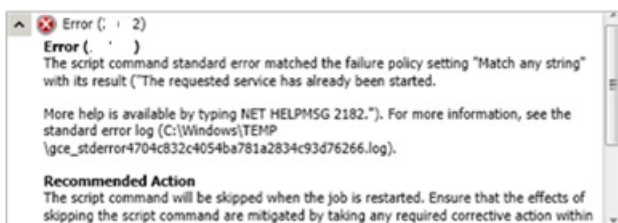
3. Select the **Timeout** seconds for this command.

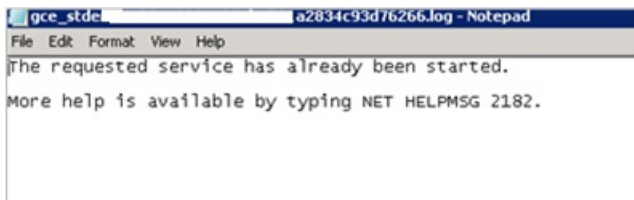
4. Click **Advanced..** and specify the output file and log file location for any errors.

5. Click **Finish**.

Monitor the job in the VMM console for the result. In case of any error, go through the error file log and follow the recommendations as applicable.

Here is a sample error message



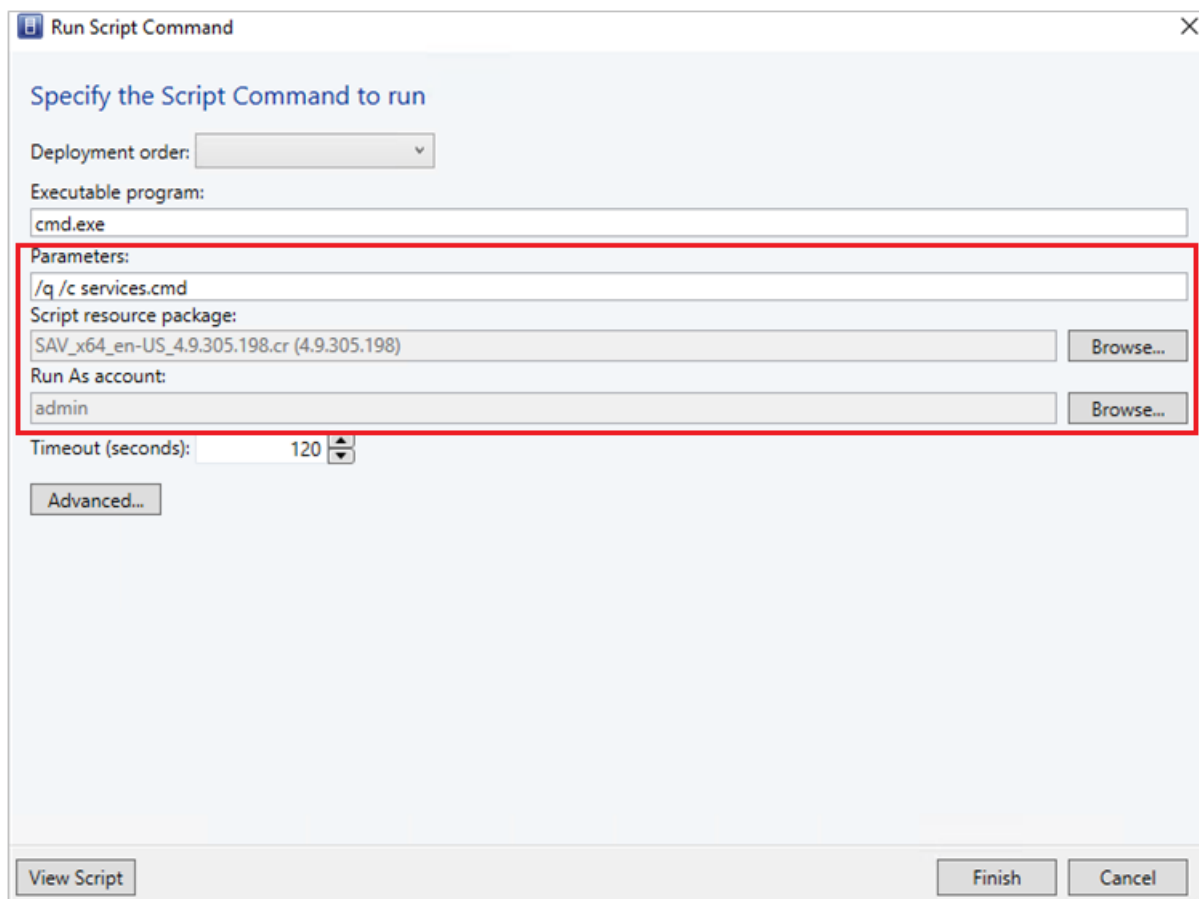


Example scenario 2 - Start a service using custom resources

In this example, at step 2, we use the following values, rest of the steps remain the same.

- **Deployment Order:** Select the deployment order from the drop-down menu.
- **Executable program:** cmd.exe
- **Parameters:** /q /c services.cmd
- **Script resource package:** SAV_x64_en-US-4.9.305.198.cr(4.9.305.198)
- **Run As account:** Admin

For this example, we included a batch file named **services.cmd** under a custom resource folder named **SAV_x64_en-US-4.9.305.198.cr(4.9.305.198)**, and specified the **Run As account** as Admin.



The batch file performs a net stop sftlist/y and then a net start sftlist. In this scenario, the custom resource folder is transferred to the agent host and copied under windows\temp. A folder with the format, **scvmm.xxxxxxxx**, is created to contain all the files. From here, it executes the batch file and the agent returns the corresponding outcome to VMM and displays it in the job. If the script generates an error, it creates a log under the specified location.

Provision a cluster from Hyper-V standalone hosts in the VMM fabric

12/7/2018 • 6 minutes to read

Use the instructions in this article to create a cluster from standalone Hyper-V host servers that are managed in the System Center - Virtual Machine Manager (VMM) fabric.

Before you start

PREREQUISITE	DETAILS
VMM	You'll need a VMM host group set up in the fabric. This is needed to allocate shared storage logical units if VMM needs to assign shared storage to the cluster nodes.
Hyper-V	<p>You must have two or more standalone Hyper-V hosts in the VMM fabric that are in the same VMM host group.</p> <p>The hosts must meet the requirements for failover clustering</p> <p>All of the hosts that will be in the cluster should be running the same operating system.</p> <p>All hosts must belong to the same VMM host group</p> <p>You must have a domain account (to use as the basis for a Run As account) for creating the cluster. The account must have administrative permissions on the servers that will become cluster nodes, and must belong to the same domain as those servers. Also, the account requires Create Computer objects permission in the container that is used for Computer accounts in the domain.</p>

PREREQUISITE	DETAILS
<p>Storage</p>	<p>Storage must be discovered and classified in the Fabric workspace of the VMM console. Then, either storage pools, logical units, or both must be allocated to the host group or the parent host group chosen for your set of hosts.</p> <p>If the shared storage isn't managed by VMM, disks must be available to all nodes in the cluster before you can add them. You'll need to provision one or more logical units to all hosts that you want to cluster, and mount and format the storage disks on one of the hosts.</p> <p>To access shared storage, the Multipath I/O (MPIO) feature must be installed on each Hyper-V host. VMM doesn't add this automatically. You can add MPIO using server manager. If MPIO is installed VMM will automatically enable it for supported storage arrays by using the Microsoft provided Device Specific Module (DSM). If you already installed vendor-specific DSMs for supported storage arrays and then add the host VMM, the vendor-specific MPIO settings will be used to communicate with those arrays. If you add a host to VMM management before you add the MPIO feature, you must add the MPIO feature, and then manually configure MPIO to add the discovered device hardware IDs. Or, you can install vendor-specific DSMs.</p> <p>If you are using iSCSI SAN as your shared storage, the Microsoft iSCSI initiator service should be installed and running (set to automatic) on each Hyper-V host. VMM uses the iSCSI initiator service to configure shared storage on the Hyper-V nodes automatically when the cluster is created. There is no need to discover iSCSI portals on each Hyper-V node if VMM manages the shared storage.</p> <p>If you are using a Fibre Channel storage array network (SAN), each host must have a host bus adapter (HBA) installed, and zoning must be correctly configured. For more information, see your storage array vendor's documentation.</p> <p>By default, when VMM manages the assignment of logical units, VMM creates one storage group per host, either a stand-alone host or a host cluster node. However, for some storage arrays, it is preferable to use one storage group for the entire cluster, where host initiators for all cluster nodes are contained in a single storage group. To support this you must set the CreateStorageGroupsPerCluster property to \$true by using the Set-SCStorageArray cmdlet.</p>

PREREQUISITE	DETAILS
Networking	<p>For all Hyper-V hosts that you want to cluster, if the hosts are configured to use static IP addresses on a particular network, make sure that the static IP addresses on all hosts are in the same subnet.</p> <p>If you've already created a network configuration in VMM that is relevant to the cluster, and have applied that configuration to network adapters in the hosts, ensure that the configuration is applied consistently across all the hosts you want to cluster. For example, if you have designated a specific set of network adapters (one per host) as management adapters for the cluster, make sure that the name of the logical network and VM network associated with those network adapters is consistent. When VMM is identifying networks that the cluster can use, it will only recognize networks with consistent settings on every node.</p>

Create a cluster

1. In the VMM console, click **Fabric > Create > Hyper-V Cluster** to open the Create Hyper-V Cluster wizard.
2. In **General**, specify a cluster name and choose the host group in which the existing Hyper-V hosts are located.
3. In **Resource Type**, select the Run As account that you'll use to create the cluster. The account that you use must have administrative permissions on the servers that will become cluster nodes, and must belong to the same domain as the Hyper-V hosts that you want to cluster. Also, the account requires **Create Computer objects** permission in the container that is used for Computer accounts in the domain. Ensure that **Existing Windows servers** is selected, and if you don't need support from Microsoft for this cluster, you can select **Skip cluster validation**.
4. In **Nodes**, select the Hyper-V host servers that you want to include in the cluster. You can select multiple hosts by using the CTRL key, or a range by using SHIFT.
5. In **IP address** (if it appears), type in the IP address you want to use for the cluster.
6. In **Storage**, select the data disks you want the cluster to use. The list of available disks includes the logical units associated with the host group that you selected at the beginning of the wizard.
 - If you assigned storage out-of-band, disks that are not managed by VMM are displayed and selected as available disks, with the check box next to each disk dimmed and unavailable.
 - If you are using a third-party clustered file system (CFS) solution, make sure you are aware which disks are CFS disks. Do not select those disks for the cluster. If you do, cluster creation will fail. If you are using a third-party clustered file system (CFS) solution, make sure you are aware which disks are CFS disks. Do not select those disks for the cluster. If you do, cluster creation will fail.
 - If the number of selected hosts for the cluster is even, the smallest disk that is larger than 500 megabytes (MB) is automatically chosen as the witness disk and is unavailable for selection.
7. In **Virtual Switches**, you can select the logical networks to use when VMM automatically create virtual switches on the Hyper-V nodes. The external virtual switches on destination Hyper-V nodes. VMM will automatically create the virtual switches on all the Hyper-V nodes.
8. In **Summary**, confirm the settings and then click **Finish**. You can monitor the cluster status on the **Jobs** page. After the job finishes you can verify cluster information by right-clicking **Properties > Status** tab on the cluster. You can also right-click the cluster and click **Validate Cluster**.

Here's what VMM does after you create the cluster:

1. Validates that all hosts meet the prerequisites, such as required operating system and domain membership
2. Enables the Failover Clustering feature on each host
3. Unmasks the selected storage logical units to each host

4. Runs the cluster validation process
5. Creates the cluster with quorum settings, configures any cluster static IP settings that you specified, and enables Cluster Shared Volumes (CSV)
6. For each logical unit that is designated as a CSV, assigns the logical unit as a CSV on the cluster

Next steps

[Provision VMs](#)

Provision a Hyper-V host or cluster from bare metal computers

3/14/2019 • 22 minutes to read

Use this article to provision a Hyper-V host or cluster from bare metal computers with nothing installed on them, in the System Center - Virtual Machine Manager (VMM) fabric.

In addition to adding existing Windows servers to the fabric as Hyper-V hosts and clusters, VMM can discover physical bare-metal machines, automatically install an operating system, and provision them as Hyper-V server hosts and clusters.

Here's how you do this:

1. **Verify prerequisites**-Make sure you have all the prerequisites before you start.
2. **Initial configuration**-Set up the BIOS on the machine to support virtualization, setting the BIOS boot order to boot from a Pre-Boot Execution Environment (PXE)-enabled network adapter as the first device, and configuring the logon credentials and IP address settings for the BMC on each computer. You need to create DNS entries and Active Directory account for the machine names, and we recommend you allow time for DNS replication to occur.
3. **Prepare the PXE server environment**: Add the PXE server to VMM management, as described in Prerequisites: creating hosts, host clusters, or Scale-Out File Server clusters from bare metal in VMM and How to add a PXE server to VMM.
4. **Add resources to VMM library**: Add resources that include a generalized virtual hard disk with an appropriate operating system (as listed in Prerequisites: creating hosts, host clusters, or Scale-Out File Server clusters from bare metal in VMM) to use as the base image, and optional driver files to add to the operating system during installation.
5. **Create profiles**: In the library, create one or more physical computer profiles. These profiles include configuration settings, such as the location of the operating system image, and hardware and operating system configuration settings.
6. **Create Hyper-V host or cluster**: You run different wizards depending on whether you want to set up a standalone host or cluster.

Before you start

NOTE

VMM doesn't support bare metal provisioning of physical Machines in disjoint namespace. As a workaround, follow these steps:

1. provision the bare metal in a non-disjoint namespace domain.
2. remove the provisioned host from VMM.
3. Join the host to the disjoint namespace domain of interest.
4. Add the host back to VMM. Use [this procedure](#).

Ensure the following prerequisites:

COMPONENT	PREREQUISITE	DETAILS
Physical computer	Support for discovery	<p>Each physical computer must have a baseboard management controller (BMC) installed that enables out-of-band management. Through a BMC, you can access the computer remotely, independent of the operating system, and control system functions such as the ability to turn the computer off or on. BMC requirements:</p> <p>The BMCs must use one of the supported out-of-band management protocols, and the management protocol must be enabled in the BMC settings.</p> <p>The supported protocols are: Intelligent Platform Management Interface (IPMI) versions 1.5 or 2.0; Data Center Management Interface (DCMI) version 1.0; System Management Architecture for Server Hardware (SMASH) version 1.0 over WS-Management (WS-Man); custom protocols such as Integrated Lights-Out (iLO)</p> <p>The BMCs should use the latest version of firmware for the BMC model.</p> <p>The BMCs must be configured with logon credentials and must use either static IP addressing or DHCP. If you use DHCP, we recommend that you configure DHCP to assign a constant IP address to each BMC, for example by using DHCP reservations.</p> <p>The VMM management server must be able to access the network segment on which the BMCs are configured.</p>
Physical computer	Hyper-V role requirements	<p>Computer that supports the Hyper-V role, must use x64-based processors and have the appropriate basic input/output system (BIOS) settings enabled.</p>
Physical computers	DNS	<p>If your environment has multiple Domain Name System (DNS) servers, where DNS replication may take some time, we strongly recommend that you create DNS entries for the computer names that are assigned to the physical computers, and allow time for DNS replication to occur. Otherwise, deployment of the computers may fail.</p>

COMPONENT	PREREQUISITE	DETAILS
Physical computer	BIOS/EFI	Determine whether the computers use Extensible Firmware Interface (EFI) or BIOS. If you have computers of each type, you must create a separate profile for each type.
Physical computers	Operating system	You can add a Windows Server 2016 node to a Windows Server 2012 R2 cluster, subject to the requirements specified previously; however, you cannot add a Windows Server 2012 R2 node to a Windows Server 2016 cluster.
PXE server	Deployment requirements	<p>You must have a PXE server configured with Windows Deployment Services.</p> <p>If you have an existing PXE server in your environment configured with Windows Deployment Services, you can add that server to VMM. Then you can use it for provisioning in VMM (and VMM recognizes only the resulting servers). All other requests continues to be handled by the PXE server according to how it is configured.</p> <p>If you don't have an existing PXE server, you can deploy the Windows Deployment Services role on a server running a supported operating system (Windows Server 2012 R2 or later).</p> <p>When you install Windows Deployment Services you should install both the Deployment server and Transport server options. Note that you don't need to add images. During host deployment, VMM uses a virtual hard disk that you've created and stored in the library. In addition you don't need to configure settings on the PXE response tab. VMM provides its own PXE provider.</p> <p>The PXE server must be in the same subnet as the physical computers that you want to provision.</p> <p>When you add a PXE server, you must specify account credentials for an account that has local administrator permissions on the PXE server. You can enter a user name and password or specify a Run As account. If you want to use a Run As account, you can create the RunAs account before you begin, or during deployment.</p>

COMPONENT	PREREQUISITE	DETAILS
PXE server	Boot order	On each computer, set the BIOS boot order to boot from a Pre-Boot Execution Environment (PXE)-enabled network adapter as the first device.
Virtual hard disk	Operating system	<p>Make sure you have a generalized virtual hard disk in a VMM library share. It should be running Windows Server 2012 R2 or later.</p> <p>We recommend that for production servers, you use a fixed disk (.vhd or .vhdx file format) to increase performance and to help protect user data. Note that by default, when you create a physical computer profile, VMM converts a dynamic disk to a fixed disk.</p> <p>If you plan to assign customer drivers they must exist in the library.</p> <p>To create the virtual hard disk, you can create a virtual machine, install the guest operating system, and then use sysprep with the /generalize and the /oobe options.</p> <p>The operating system on the virtual hard disk that you deploy on hosts or clusters must support the boot from virtual hard disk (VHD) option.</p> <p>If you use Remote Desktop Services (RDS) to manage servers, we recommend that you enable the RDS connections in the image. You can also enable RDS by using an answer file in the physical computer profile.</p>

COMPONENT	PREREQUISITE	DETAILS
Networking	Logical networks	<p>If you have already configured logical networks or logical switches in VMM, you can include those configurations in a physical computer profile.</p> <p>To include a logical switch that you want to apply to physical NICs in a physical computer profile (for hosts or host clusters), you must first take certain steps. Make sure that you have installed the intended number of NICs on the host computer or computers. In addition, before you create the physical computer profile in VMM, create the logical switch.</p> <p>To include static IP addressing controlled through a logical network in a physical computer profile configure the logical network. The logical network must include at least one network site and static IP address pool. The network site must also be available to the host group or to a parent host group where you want to assign the hosts that you are creating from bare metal.</p> <p>To include a virtual NIC for hosts in a physical computer profile (for hosts or host clusters), you must first take certain steps. Make sure that you have installed the intended number of physical NICs on the computer or computers that are to become hosts. In addition, on the VMM management server, install all necessary virtual switch extensions and extension providers, create at least one VM network, and create a logical switch. In the logical switch, as a best practice, include one or more port classifications for the virtual ports.</p>

COMPONENT	PREREQUISITE	DETAILS
Physical computer profile	Answer file	<p>If you want a physical computer profile to include references to an answer file (Unattend.xml file) or to custom resources (for example, an application installer that is referenced in post-deployment script commands), create the answer file or obtain the custom resources before deployment, and add them to a VMM library share. Within a library share, place custom resources in one or more folders with a .CR (custom resource) extension. VMM recognizes them as custom resources. For example, you might want to create an answer file to enable Remote Desktop Services and place it on a library share. Then you can select that file when you configure a physical computer profile.</p> <p>By default when you deploy servers or clusters from bare metal, VMM automatically performs the following (no answer file or post-deployment commands needed): Installs the Hyper-V role for Hyper-V hosts. Installs the Hyper-V role, failover cluster feature, and multipath I/O (MPIO) for Hyper-V clusters.</p>
Accounts	<p>You need two Run As accounts.</p> <p>A Run As account for joining computers to the domain. You can create a Run As account in the Settings workspace.</p> <p>A Run As account for access to the baseboard management controller (BMC) on each computer.</p>	

Add a PXE server to the VMM fabric

1. Click **Fabric > Servers > Add > Add Resources > PXE Server**.
2. In Computer name specify the PXE server name.
3. Add the credentials for an account that has local administrator permissions on the PXE server. You can specify an existing Run As account (or click Create Run As Account to create a new one), manually enter user credentials in the format domain_name\user_name. Click **Add**.
4. In **Jobs** verify that the job status is Completed and close the dialog box. The job sets up the new PXE server, installs the VMM agent on the PXE server, imports a new Windows Preinstallation Environment (Windows PE) image, and adds the machine account for the PXE server to VMM.
5. Verify that the PXE server is added in **Fabric > Servers > PXE Servers > Home > Show > Fabric Resources > PXE Servers**. The agent status should be **Responding**.

Add and assign driver files

If you plan to assign custom drivers, the driver files must exist in the library. You can tag the drivers in the library, so that you can later filter them by tag. After files are added, when you configure a physical computer profile, you

can specify the driver files. VMM installs the specified drivers when it installs the operating system on a physical computer.

In the physical computer profile, you can select to filter the drivers by tags, or you can select to filter drivers with matching Plug and Play (PnP) IDs on the physical computer. If you select to filter the drivers by tags, VMM determines the drivers to apply by matching the tags that you assign to the drivers in the library to the tags that you assign in the profile. If you select to filter drivers with matching PnP IDs, you don't need to assign custom tags.

1. Locate a driver package that you want to add to the library.
2. In the library share that is located on the library server that is associated with the group where you want to deploy the physical computers, create a folder to store the drivers, and then copy the driver package to the folder.
3. We strongly recommend that you create a separate folder for each driver package, and that you do not mix resources in the driver folders. If you include other library resources such as .iso images, .vhd files or scripts with an .inf file name extension in the same folder, the VMM library server is unable to discover the resources. Also, when you delete an .inf driver package from the library, VMM deletes the entire folder where the driver .inf file resides.
4. In the VMM console, open the Library workspace. In the **Library > Library Servers**, expand the library server where the share is located, right-click the share, and then click **Refresh**. After the library refreshes, the folder that you created to store the drivers appears.
5. Now assign tags if required. In **Library**, expand the folder that you created to store the drivers in the previous procedure, and then click the folder that contains the driver package.
6. In the **Physical Library Objects**, right-click the driver .inf file, and then click **Properties**.
7. In the **Driver File Name Properties > Custom tags**, enter custom tags separated by a semi-colon, or click **Select** to assign available tags or to create and assign new ones. If you click **Select**, and then click **New Tag**, you can change the name of the tag after you click **OK**. For example, if you added a network adapter driver file, you could create a tag that is named ServerModel NetworkAdapterModel, where ServerModel is the server model and NetworkAdapterModel is the network adapter model.

Create a physical computer profile

1. Click **Library > Home > Create > Physical Computer Profile**.
2. In the **New Physical Computer Profiles Wizard > Profile Description**, type in a name and description and select **VM host**.
3. In **OS Image > Virtual hard disk file > Browse**, click the generalized virtual hard disk that you added to the library share. By default, if the disk is dynamic VMM converts it to a fixed disk during host deployment. We recommend that for production servers you use a fixed disk to increase performance and help protect user data.
4. In **Hardware Configuration > Management NIC** select the network adapter used to communicate with VMM and whether to use DHCP or a static address. If you want to use Consistent Device Naming (CDN) for the adapter or configure logical switches and ports click **Physical Properties**. Click **Add** to add the adapter.
5. In **Disk**, specify the partitioning scheme for the first disk. You can use GPT if the physical computer profile is EFI. In **Partition Information**, select the volume label, whether to use all remaining free space or a specific size, and whether to designate the partition as the boot partition. You can also add a new disk or partition. During deployment VMM copies the virtual hard disk file to the boot partition and automatically create a system partition on the same disk.
6. In **Driver filter**, filter the driver filters to be applied to the operating system during host deployment. You can filter by Plug and Play ID or by specific tags. If you select to filter drivers with matching tags make sure you've added driver files to the library and assigned corresponding tags.
7. In **OS Configuration** specify the domain that the Hyper-V host or cluster should join, specify local admin credentials and identity information. Add the product key for installation, and set the time zone. In

GUIRunOnce, specify one or more commands to run when the user logs on to the Hyper-V host for the first time.

8. In **Host Settings** specify the path of the host to store the files that are associated with virtual machines placed on the host. Don't specify drive C because it's not available for placement. If you don't specify a path, VMM placement determines the most suitable location.
9. In **Summary** verify the settings. Wait until Jobs shows a status of completed, and then verify the profile in **Library > Profiles > Physical Computer Profiles**.

PCP post deployment settings

After you successfully create and deploy the PCP, you can configure additional settings such as RDMA, QoS, and SET using the PCP post deployment script.

Sample script

Here is the sample script to configure RDMA, SET and QoS.

```
# Install data center bridging
Install-WindowsFeature Data-Center-Bridging

#Enable RDMA, assuming customer chosen switch name for storage as Storage1Switch and Storage2Switch
Enable-NetAdapterRDMA "Storage1Switch"
Enable-NetAdapterRDMA "Storage2Switch"

# set Qos Policy
New-NetQosPolicy "SMB" -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3

# Enable net qos flow control
Enable-NetQosFlowControl -Priority 3

# Disable net qos flow control other than 3
Disable-NetQosFlowControl -Priority 0,1,2,4,5,6,7

# Enable net adapter qos on all adapters
Enable-NetAdapterQos -InterfaceAlias "*"

# set qos traffic class
New-NetQosTrafficClass "SMB" -Priority 3 -BandwidthPercentage 50 -Algorithm ETS

# Install windows feature
Install-WindowsFeature -Name Hyper-V

Install-WindowsFeature -Name RSAT-Hyper-V-Tools

# set net adapter property "encapsulated overhead"
NetAdapterAdvancedProperty -Name "*" -DisplayName "Encapsulated Overhead" -DisplayValue "160"

#disable ipv6
netsh int ipv6 isatap set state disabled

#Configure SET team mapping between virtual network adapter to physical network adapters. (Note: to get names
of adapters, use command Get-NetAdapater. For team mapping use command
$physicalAdapters = Get-NetAdapter -Physical
$virtualStorageAdapter1 = Get-VMNetworkAdapter -ManagementOS | Where-Object {$_.Name -eq "Storage1Switch"}
$virtualStorageAdapter1 = Get-VMNetworkAdapter -ManagementOS | Where-Object {$_.Name -eq "Storage2Switch"}

Set-VMNetworkAdapterTeamMapping -ManagementOS -PhysicalNetAdapterName $physicalAdapters[0].Name -
VMNetworkAdapterName $virtualStorageAdapter1.Name
Set-VMNetworkAdapterTeamMapping -ManagementOS -PhysicalNetAdapterName $physicalAdapters[1].Name -
VMNetworkAdapterName $virtualStorageAdapter2.Name

#Set firewall rules.
[System.String[]]$Alias=@("vEthernet (WssdStorage2)", "vEthernet WssdStorage1");
$Profile ='Any'
$Name="File and Printer Sharing"
$rules = Get-NetFirewallRule -DisplayGroup $Name
```

```

foreach ($rule in $rules)
{
    $rule | Get-NetFirewallAddressFilter | Set-NetFirewallAddressFilter -
        LocalAddress Any -RemoteAddress Any
}

Set-NetFirewallRule -DisplayGroup $Name -Enabled True -Profile $Profile - InterfaceAlias $Alias
$Profile='Any'
$Name=='FPS-LLMNR-In-UDP'
Set-NetFirewallRule -Name $Name -Enabled True -Profile $Profile
[System.String[]]$Alias=@("Storage2Switch", "Storage1Switch", "ManagementSwitch");
$Profile ='Any'
$Name="Windows Remote Management"
Set-NetFirewallRule -DisplayGroup $Name -Enabled True -Profile $Profile -InterfaceAlias $Alias

#Set assurance settings
reg add HKLM\SYSTEM\CurrentControlSet\Control\DeviceGuard /v EnableVirtualizationBasedSecurity /t REG_DWORD
/d 1 /f
reg add HKLM\SYSTEM\CurrentControlSet\Control\DeviceGuard /v RequirePlatformSecurityFeatures /t REG_DWORD /d
2 /f
reg add HKLM\SYSTEM\CurrentControlSet\Control\LSA /v LsaCfgFlags /t REG_DWORD /d 1 /f
reg add HKLM\SYSTEM\CurrentControlSet\Control\LSA /v DisableRestrictedAdmin /t REG_DWORD /d 0 /f

```

Provision a Hyper-V host from bare metal

When you deploy a Hyper-V host from bare metal, VMM does the following:

1. Discovers the physical computer through out-of-band management.
2. Deploys an operating system image on the computer through the physical computer profile.
3. Enables the Hyper-V role on the computer.
4. Brings the computer under VMM management as a managed Hyper-V host.

Provision as follows:

1. Click **Fabric > Servers > Home > Add > Add Resources > Hyper-V Hosts and Clusters**.
2. In the **Add Resource Wizard > Resource location**, select **Physical computers to be provisioned as virtual machine hosts**.
3. In **Credentials and Protocol** select the Run As account with permissions to access the BMC. In the **Protocol** list, click the out-of-band management protocol that your BMCs use. If you want to use Data Center Management Interface (DCMI), click Intelligent Platform Management Interface (IPMI). Although DCMI 1.0 is not listed, it is supported. Make sure the correct port is selected.
4. In **Discovery Scope**, specify the IP address scope that includes the IP addresses of the BMCs. You can enter a single IP address, an IP subnet, or an IP address range.
 - If you are provisioning a single computer, you can either specify a single IP address, or specify an IP address range that both starts and ends with the intended IP address. If you specify a single IP address, when you click **Next**, the computer is restarted.
 - If you specify an IP address range, when you click **Next**, information about the computer is displayed, and you can confirm that you specified the computer that you meant to.
5. If you specified a single IP address on the previous page, skip this step. Otherwise, the **Target Resources** page appears. Review the list of discovered BMCs (identified by IP addresses), and select the ones you want to provision as hosts. If you don't see all the BMCs that you expect, confirm that they are on a network accessible to the VMM server, and as needed, click **Refresh**.
6. In **Provisioning Options**, click a host group for new Hyper-V hosts. Select the physical computer profile you want to apply.
7. In **Deployment Customization**, review the list of computers again, and provide information for each computer that you want to include.

- If you see a computer that you don't want to include, select the BMC (identified by IP address) and then click **Remove**.
 - To configure computers click the BMC IP address.
 - Specify a unique computer name, without wildcard characters.
 - Select or clear **Skip Active Directory for this computer name**. The Active Directory check prevents deployment if the computer account already exists. This helps prevent deploying a computer with the same name as an existing computer. Note that if you skip the Active Directory check, and there is an existing computer account in AD DS other than the Run As account that was specified in the physical computer profile, the deployment process fails to join the computer to the domain.
8. For the computer you are configuring, click a network adapter (on the left). You can modify the configuration, or fill in more information (...).
 9. You can specify the MAC address of the management NIC (not the BMC) and static IP settings for this network adapter. If you specify an address select a logical network and an IP subnet if applicable. If the selected IP subnet includes IP address pool, you can check **Obtain an IP address corresponding to the selected subnet**. Otherwise, type an IP address that's within the logical network or its subnet. If you select an IP subnet, make sure that it corresponds to the physical location where you are deploying the host and to the network that the adapter is connected to. Otherwise, deployment may fail.
 10. Configure the adapter settings for each network adapter. Note that if the number of physical network adapters in a computer does not match the number of physical network adapters that are defined in the physical computer profile, you must specify any information that is missing for the adapters. If you decide not to provision this computer right now (for example, if physical hardware needs to be installed or uninstalled), you can select the computer's BMC IP address from the list and then click **Remove**.
 11. Repeat the configuration for each BMC IP address in the list. When you have filled in information for all the computers you want to provision, click **Next**.
 12. In **Summary**, confirm the settings, and then click **Finish** to deploy the new Hyper-V hosts and bring them under VMM management. Depending on your settings, the **Jobs** dialog box might appear. Make sure that all steps in the job have a status of **Completed**, and then close the dialog box.
 13. To confirm that the host was added click **Fabric > Servers > All Hosts > host group**, and verify that the new Hyper-V host appears in the group.

Provision a Hyper-V cluster from bare metal

When you deploy a Hyper-V cluster from bare metal VMM does the following:

1. Discovers the physical computers through out-of-band management.
2. Deploys an operating system image on the computers by using the selected physical computer profile.
3. Installs the failover clustering feature, and the Hyper-V role and MPIO feature.
4. Brings the provisioned cluster under VMM management

Provision as follows:

1. Click **Fabric > Servers > Add > Add Resources > Hyper-V Hosts and Clusters**.
2. In **General Configuration**, specify a name for the host cluster. Choose a storage configuration if required:
 - For shared storage, click **Storage connected to the cluster using shared SAS, FC, or iSCSI**.
 - For Storage Spaces Direct, click **Disk subsystem directly connected to individual nodes in the cluster**.
3. In **Resource Type >** select **Physical computers to be provisioned**:
 - Specify the administrator Run As account to use for creating the cluster.
 - Select the physical computer profile (which provides the domain name and administrator Run As account for each node).
 - Next to the **BMC Run As** account box, click **Browse**, and select a Run As account that has permissions

to access the BMC.

- In the **Out-of-band management** protocol list, click the protocol that your BMCs use. If you want to use Data Center Management Interface (DCMI), click Intelligent Platform Management Interface (IPMI). Although DCMI 1.0 is not listed, it is supported. Ensure that the correct port is selected.
 - If the **Skip cluster validation** option appears, and you don't need support from Microsoft for this cluster, you can skip validation.
4. In **Discovery Scope**, specify the IP address scope that includes the IP addresses of the BMCs. You can enter a single IP address, an IP subnet, or an IP address range. Deep discovery provides detailed information about a computer (for example, MAC addresses of network adapters) but restarts the computer, and requires additional time. You can allow or skip deep discovery.
 5. If you specified a single IP address on the previous page, skip this step. Otherwise, the **Target Resources** page appears. Review the list of discovered BMCs (identified by IP addresses), and select the ones you want to include in the cluster.
 6. If you don't see all the BMCs that you expect, confirm that they are on a network accessible to the VMM server, and as needed, click **Refresh**. Allow or skip deep discovery. Deep discovery provides detailed information about a computer (for example, MAC addresses of network adapters) but restarts the computer, and requires additional time. Then click **Next**.
 7. In **Deployment Customization** provide information for each computer that you want to include. If you see a computer that you don't want to include, select the BMC (identified by IP address) and then click **Remove**.
 - To configure computers click the BMC IP address. Specify a unique computer name, without wildcard characters.
 - Select or clear **Skip Active Directory for this computer name**. The Active Directory check prevents deployment if the computer account already exists. Note that if you skip the check and there's an existing computer account in AD other than the Run As account that was specified in the physical computer profile, the deployment process fails to join the computer to the domain.
 - For the computer you are configuring, click a network adapter. You can modify the configuration, or fill in more information.
 - You can specify the MAC address of the management NIC (not the BMC) and static IP settings for this network adapter. If you specify an address select a logical network and an IP subnet if applicable. If the selected IP subnet includes IP address pool, you can check **Obtain an IP address corresponding to the selected subnet**. Otherwise, type an IP address that is in within the logical network or its subnet. If you select an IP subnet, make sure that it corresponds to the physical location where you are deploying the host and to the network that the adapter is connected to. Otherwise, deployment may fail.
 8. Configure the network adapter settings for each network adapter. Note that if the number of physical network adapters in a computer does not match the number of physical network adapters that are defined in the physical computer profile, you must specify any information that is missing for the adapters. If you decide not to provision this computer right now (for example, if physical hardware needs to be installed or uninstalled), you can select the computer's BMC IP address from the list and then click **Remove**.
 9. Repeat the configuration for each BMC IP address in the list.
 10. When you have filled in needed information for all the computers you want to provision, click Next.
 11. In **Summary** confirm the settings, and then click **Finish** to deploy the new Hyper-V hosts and bring them under VMM management. Depending on your settings, the Jobs dialog box might appear. Make sure that all steps in the job have a status of Completed, and then close the dialog box.
 12. To confirm that the host was added click **Fabric > Servers > All Hosts >** and locate and click the new host cluster. In the **Hosts** pane, in the **Host Status** column, verify that each node in the cluster is OK.

Create a guest cluster from a VMM service template

6/14/2019 • 8 minutes to read

Use this article if you want to create a guest failover cluster using a System Center - Virtual Machine Manager (VMM) service template.

A guest failover cluster consists of multiple VMs that are deployed in a cluster and use shared storage. Services in VMM are used to group together virtual machines to provide an app. Service templates contain information about a service, including the VMs that are deployed as part of the service, the applications to install on VMs, and the network configuration that should be used. You can add VM templates, network settings, applications and storage to a service template. [Learn more](#)

You can use service templates to create a guest cluster . That cluster can then be configured to run an app such as SQL Server.

Before you start

- VMs in a guest cluster can only be deployed to Hyper-V host clusters running Windows Server 2012 R2 or later. Otherwise deployment will fail.
- You can deploy a guest failover cluster that uses shared .vhdx files on a Hyper-V failover cluster. In this scenario, if Hyper-V uses Cluster Shared Volumes (CSV) on block-level storage, then the shared vhdx files are stored on a CSV that's configured as shared storage.
- Alternatively, Hyper-V can use SMB file-based storage deployed by Scale-Out File Server (SOFS), as the location of the shared .vhdx files.
- No other shared storage types are supported for guest clusters. Third-party SMB storage isn't supported.
- You need a number of scripts to create the guest cluster, including a script to run on the first VM in the cluster, a script to run on the other VMs so that they can join the cluster. Script settings are specified in the service template application settings.
- To configure shared disks for the cluster you'll need to use new VHDX files. Don't reuse from a previous cluster. Ensure that the hard disk files are in the VMM library.
- Identify a single path in SCSI-based storage where all the VHDX files for the guest cluster will be placed at deployment time. You can use storage classifications to control the placement of VHDX files but you'll need at least one location in the classification with the capacity to hold all of the VHDX files. VMM doesn't deploy the VHDX files to multiple locations.
- You can vary the location of VHDX files at deployment time, even if you use the same service template to deploy multiple guest clusters. To do this you'll need to deploy the guest clusters to a host group and not a cloud. Then at deployment you specify a single path for all the shared VHDX files for the cluster. This overrides the location specified in the VM template.
- You'll need a virtual hard disk file that contains the operating system (prepared with Sysprep) that you want the VMs in the guest cluster to use. When each node is created VMM uses a copy of the virtual hard disk file for the system disk of the node.

Specify scripts that run when a guest cluster is created

1. [Set up an application profile.](#)
2. In **New Application Profile > General > Compatibility** leave the default **General** setting enabled.
3. In **Application Configuration > OS Compatibility** select one or more editions of a server operating system (not earlier than Windows Server 2012).

4. Add the scripts you need for creating the first node of the cluster and then adding other nodes. Provide the scripts as follows:
 - For a script that will run on the first node of the cluster when it is created (and not on other nodes), for **Script command type**, select **Creation: First VM**.
 - For a script that will run on later nodes of the cluster when they are created (and not on the first node), for **Script command type**, select **Creation: VMs After First**.
 - For each script, specify the executable name and the parameters through which the script will run, and the RunAs account.
 - Configure other settings as needed, including how long the script should run before timing out, failure and restart policies.
 - A script can contain settings to be entered when you are configuring the service for deployment. To format this type of setting, type the parameter in the **Parameters** field in the following format:
@<SettingLabel>@ (for example, type @ClusterName@).
 - Example: a script FormCluster.exe that runs with Cmd.exe and the /q and /c parameters and requires the cluster name would be Executable Program: **Cmd.exe**, Parameters: **/q /c FormCluster.cmd @ClusterName@**
 - You can also add scripts to delete the cluster in an orderly way. **Script command type** would be **Deletion: VMs Before Last** or **Deletion: Last VM**.
 - You can also add a script of type **Pre-Install** that will run on the first VM and on later VMs that are created as part of the service tier.
5. Click **OK** to save settings and verify that the profile was created in **Profiles > Application Profiles**. The profile will appear in the **Profiles** pane.

Create a VM template

Create a VM template that includes settings for a shared VHDX file. The VHDX file must be deployed to shared storage that has SCSI channels available for each cluster node to provide the same access to the file for each node.

1. In the VMM library, verify that you have a virtual hard disk that contains the operating system (created using SysPrep) you want to use for VM in the guest cluster. It mustn't be blank.
2. Create a VM template.
3. In the **Create VM Template Wizard > Select Source**, select **Use an existing VM template or virtual hard disk stored in the library > Browse**.
4. In **Select VM Template Source**, click the virtual hard disk you want to use.
5. In **Configure Hardware**, specify a hardware profile or hardware settings.
6. To configure the guest cluster to use a shared VHDX, in **Bus Configuration** click **SCSI adapter 0**, and then next to **New** click **Disk**. The new disk appears as a listing under the SCSI adapter. Select the disk and select **Share the disk across the service tier**.
7. Clear **Contains the operating system for the virtual machine**. Click **Browse**, and select the VHDX file that you want VMM to deploy to shared storage and click **OK**. Repeat for each additional node in the cluster. Add the same disk each time but ensure that the SCSI channel is unique for each node.
8. In **Network Adapters** select the adapter, and select **Enable guest specified IP addresses**. This enables the nodes (VMs) in the cluster to specify IP addresses for the cluster itself and for applications that you configure to run in the cluster.
9. In **Advanced > Availability**, select **Make this virtual machine highly available**. With this setting enable the VM is created as a clustered instance on the host cluster so that if one host fails the VM will fail over to another host in the cluster.
10. Click **Manage availability sets > Create**. The availability set you create will be used by all the nodes in the guest cluster. This means that VMM will attempt to keep the VMs on separate hosts, so that if one host fails VMs on a different host can provide services.

11. In **Configure Operating System**, open the **Guest OS profile** list and select a guest operating system profile or **Create new Windows operating system customization settings**. Your selection determines whether additional wizard pages are displayed.
 - Under **Identity Information** > **Computer name**, you can provide a pattern to generate computer names. For example, if you enter `server####`, the computer names that are created are `server0001`, `server0002`, and so on. The use of a pattern ensures that when you add additional virtual machines to a service, the computer names that are generated are related and identifiable. If you use this method to specify the computer name, you cannot use it in combination with a name prompt parameter (`@<name>@`). You can use one method or the other, but not both.
 - Under **Networking**, you can specify Active Directory settings by using the FQDN or by using at signs (@) before and after the domain name, for example, `@Domain@`. By using the at signs (@) in this way, the necessary information can be entered when the virtual machine is deployed as part of a service. You don't need a trust relationship between the domain in which the service is deployed and the VMM management server domain.
12. Finish the wizard to create the VM template.

Include the VM template in a service template

1. [Create a service template](#), and add the VM template to the appropriate template tier.
2. After you save and validate the template right-click the tier object in the service template designer and click **Properties**.
3. In **Application Configuration**, add the application profile you created. When the service is deployed the scripts in the application profile will run. Save and validate the service template.
4. Right-click the service template again > **Properties**.
5. In **General**, select **This machine tier can be scaled out** and specify a value greater than 1 for **Default instance count** and **Maximum instance count**. Maximum should be set to less or equal to the number of SCSI channels you configured in the VM template. The default count should be less than the maximum.
6. In **Number of upgrade domains**, specify the same value as that in **Maximum instance count**. For example, if you specify a default count of 3 and a maximum count of 3 the guest cluster will have three nodes. The number of upgrade domains should also be set to 3 so that updates are performed in three stages, one node (VM) at a time. This leaves at least two VMs in the guest cluster running during planned maintenance.
7. Save and validate the service template.

After you've set up the guest cluster you're ready to deploy the service.

Next steps

[Deploy VMs from a template](#)

Set up networking for Hyper-V hosts and clusters in the VMM fabric

5/30/2019 • 4 minutes to read

Read this article to set up network settings for Hyper-V hosts and clusters in the System Center - Virtual Machine Manager (VMM) compute fabric.

You can apply network settings to a Hyper-V host or cluster using a logical switch. Applying a logical switch ensure that logical networks, and other network settings, are consistently assigned to multiple physical network adapters.

Before you start

- If you want to configure network settings manually ensure you've set up logical networks before you begin. In addition make sure that the network sites within your logical networks are configured to use the host group of the host you want to assign them to. Check this in **Fabric > Servers > All Hosts**, and click the host group. In **Hosts**, click the host > **Properties**.
- If you want to use a logical switch you need to create the logical switch and port profiles.

Configure network settings with a logical switch

To do this you'll need to [configure the logical switch](#) and [port profiles](#) you'll apply. Then you need to indicate what the physical network adapter is used for, and configure network settings by applying a logical switch. The network adapters that you configure can be physical or virtual adapters on the hosts.

Specify what the network adapter is used for

Regardless of any port profiles and logical switches you are using in your network configuration, you must specify whether a network adapter in a host is used for virtual machines, host management, neither, or both. (The host must already be under management in VMM.)

1. Open **Fabric > Servers > All Hosts > host-group-name > Hosts > Host > Properties > Hardware**.
2. Under **Network adapters**, click the physical network adapter that you want to configure.
 - If you want to use this network adapter for virtual machines, ensure that **Available for placement** is checked.
 - If you want to use this network adapter for communication between the host and the VMM management server, ensure that **Used by management** is checked. You must make sure that you have at least one network adapter available for communication between the host and the VMM management server.
3. You don't need to configure individual settings in **Logical network connectivity** because you're using a switch.

Apply a logical switch

1. Open **Fabric > Servers > All Hosts > host group > Hosts > Host > Properties > Virtual Switches**.
2. Select the logical switch you created. Under **Adapter**, select the physical adapter that you want to apply the logical switch to.
3. In the **Uplink Port Profile** list, select the uplink port profile that you want to apply. The list contains the uplink

port profiles that have been added to the logical switch that you selected. If a profile seems to be missing, review the configuration of the logical switch and then return to this property tab. Click **OK** to finish. Note that if you didn't create the virtual switch earlier and do in now, the host might temporarily lose network connectivity when VMM creates the switch.

4. Repeat the steps as needed. If you apply the same logical switch and uplink port profile to two or more adapters, the two adapters might be teamed, depending on a setting in the logical switch. To find out if they will be teamed, open the logical switch properties, click the **Uplink** tab, and view the **Uplink mode** setting. If the setting is **Team**, the adapters will be teamed. The specific mode in which they will be teamed is determined by a setting in the uplink port profile.
5. After applying the logical switch you can check that the network adapter settings and verify whether they're in compliance with the switch:
 - Click **Fabric > Networking > Logical Switches > Home > Show > Hosts**.
 - In **Logical Switch Information for Hosts** verify the settings. **Fully compliant** indicates that the host settings are compliant with the logical switch. **Partially compliant** indicates some issues. Check the reasons in **Compliance errors**. **Non compliant** indicates that none of the IP subnets and VLANs defined for the logical network are assigned to the physical adapter. Click the switch > **Remediate** to fix this.
 - If you have a cluster, check each node.

Monitor physical network devices

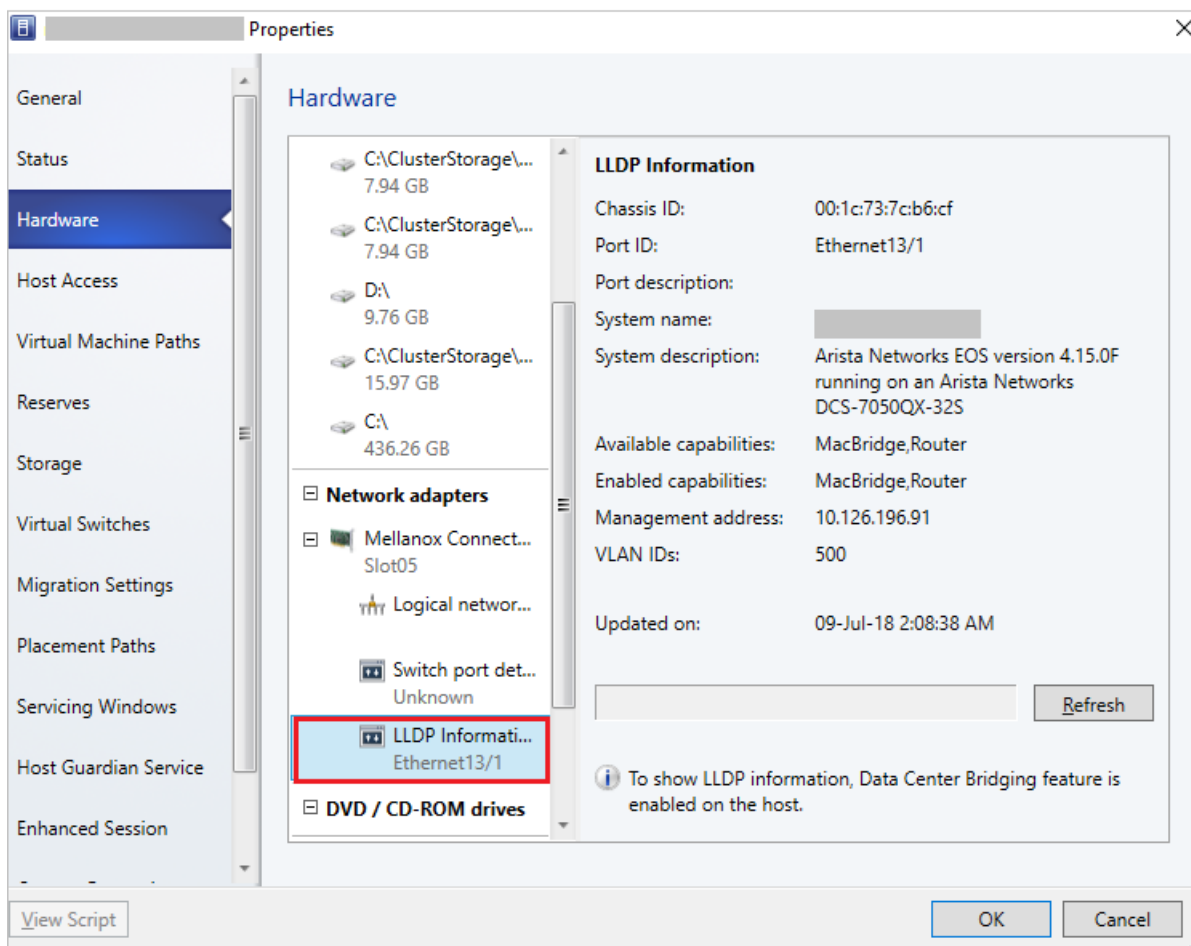
VMM supports Link Layer Discovery Protocol (LLDP). You can now leverage the LLDP information to remotely monitor physical network device properties and information. You can view this information by using VMM console and PowerShell.

Console view

To get the details of network devices from the VMM console, go to **View > Host > Properties > Hardware Configuration > Network adapter**.

NOTE

The details displayed contain a time stamp (updated on). To get the current details, refresh the page.



The following LLDP information is displayed:

INFORMATION DISPLAYED	DESCRIPTION
Chassis ID	Switch chassis ID
Port ID	Switch port to which NIC is connected
Port Description	Details related to the port such as <i>Type</i>
System Name Manufacturer	Manufacturer, Software version details
System Description	Detailed system description
Available Capabilities	Available system capabilities (such as switching, routing)
Enabled Capabilities	Enabled system capabilities (such as switching, routing)
VLAN ID	Virtual LAN identifier

INFORMATION DISPLAYED	DESCRIPTION
Management Address	IP management address

PowerShell

Use the following PowerShell command to view/refresh the LLDP details:

```
Set-SCVMHostNetworkAdapter -RefreshLLDP
```

NOTE

By default, LLDP Packet wait time is set as 30 seconds. You can modify this value by modifying the registry key at **Software\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings\LLdpPacketWaitIntervalSeconds**. The minimum value you can set is 5 seconds, maximum value is 300 seconds.

Next steps

[Set up storage](#) for Hyper-V hosts.

Add storage to Hyper-V hosts and clusters

12/7/2018 • 8 minutes to read

Read this article to allocate provisioned storage to Hyper-V hosts and clusters in the System Center - Virtual Machine Manager (VMM) fabric.

Before you start

Before you can allocate provisioned storage to hosts and cluster, it should be discovered and classified in the VMM fabric:

1. Discover and classify storage:
 - [Add and classify block storage devices](#). Learn about [classification](#).
 - [Add file storage](#)
2. [Allocate block storage to host groups](#). You can allocate an entire storage pool, or a specific logical unit (LUN).
3. Make sure you've completed these steps before you allocate storage to hosts:
 - **MPIO**: If you're using fiber channel or iSCSI storage, the Multipath I/O (MPIO) feature must be enabled on each host.
 - If MPIO is already enabled before you add the host, VMM will automatically enable it for supported storage arrays using Microsoft DSM. If you have vendor specific DSMs these will be used.
 - If you add a host to VMM and enable MPIO later, you need to configure it manually to add the discover device hardware IDs.
 - **HBA and zoning**: If you're using Fiber Channel storage array network (SAN), each host must have a host bus adapter (HBA) installed, and zoning must be correctly configured.
 - **iSCSI**: If you are using an iSCSI SAN, make sure that iSCSI portals have been added, and that the iSCSI initiator is logged into the array.

Make sure that the Microsoft iSCSI Initiator Service on each host is started and set to **Automatic**

- **Storage group**: Explain to your storage administrator how VMM manages storage.
 - In VMM a storage group binds together host initiators, target ports, and logical units.
 - A storage group contains one or more host initiator IDs (IQN or WWN) (WWN).
 - A storage group also contains one or more target ports and one or more logical units. Logical units are exposed to the host initiators through the target ports.
 - By default, when VMM manages the assignment of logical units, VMM creates one storage group per host, either a standalone host or a host cluster node.
 - For some storage arrays, it is preferable to use one storage group for the entire cluster, where host initiators for all cluster nodes are contained in a single storage group. To do this, you need to set the CreateStorageGroupsPerCluster property to \$true by using the Set-SCStorageArray cmdlet.

Allocating storage

- You can allocate file storage directly to hosts and clusters.
- You can add LUNs to hosts and clusters:
- If you already provisioned LUNs on a host group, you can assign these to hosts and clusters.
- If you provisioned a storage pool on a host group, you can create LUNs during the procedure to add storage to

a cluster.

- If you want to use shared storage that isn't managed by VMM, the storage disks must be available to all hosts or nodes before you can add them. You need to provision one or more LUNs to all hosts you want to cluster, and then mount and format the storage disks on one of the nodes. Note that VMM doesn't support or block the use of asymmetric storage, where a workload can use disks that are shared between a subset of the cluster nodes. Each cluster node must be a possible owner of the cluster disk.
- After adding iSCSI storage to a host, you need to create a new session to the storage.

Allocate file storage to a standalone host

You can assign file shares on any host on which you want to create VMs that will use the file share as storage.

1. Click **Fabric > Servers > All Hosts**, and select the host or cluster node you want to configure.
2. Click **Host > Properties > Host Access**. Specify a Run As account. By default the Run As account that was used to add the host to VMM is listed. In the **Run As** account box, configure the account settings. You can't use the account that you use for the VMM service. Note that:
 - If you used a domain account for the VMM service account, add the domain account to the local Administrators group on the file server.
 - If you used the local system account for the VMM service account, add the computer account for the VMM management server to the local Administrators group on the file server. For example, for a VMM management server that is named VMMServer01, add the computer account VMMServer01\$.
 - Any host or host cluster that accesses the SMB 3.0 file share must have been added to VMM using a Run As account. VMM automatically uses this Run As account to access the SMB 3.0 file share.
 - If you specified explicit user credentials when you added a host or host cluster, you can remove the host or cluster from VMM, and then add it again by using a Run As account.
3. Click **Host Name Properties > Storage > Add File Share**.
4. In **File share path**, select the required SMB 3.0 file share, and then click **OK**.
5. To confirm that the host has access, open the **Jobs** workspace to view the job status. Or, open the host properties again, and then click the **Storage** tab. Under **File Shares**, click the SMB 3.0 file share. Verify that a green check mark appears next to **Access to file share**.
6. Repeat this procedure for any standalone host that you want to access the SMB 3.0 file share, or for all nodes in a cluster

Assign a logical unit to a standalone host

Either assign an existing unit, or create a new one and assign it.

1. In **Fabric > Servers > All Hosts**, right-click the host that you want to configure > **Properties**.
2. If you want to create a new logical unit:
 - On the toolbar, next to **Disk**, click **Add**. Next to **Logical unit** click **Create Logical Unit**.
 - In Create Logical Unit > **Storage pool** choose the pool from which to create the logical unit. Specify a name (alphanumeric only), a description and the unit size. Click **OK** to finish.
3. To assign an existing logical unit to the host, on the toolbar, next to **Disk**, click **Add**, and select the logical unit you want to assign.
4. In the **Logical unit** list, verify that the logical unit that you just created is selected.
5. In **Format new disk**, if you want to format the disk, select **Format this volume as NTFS volume with the following settings**, and specify the settings. Note that if you select **Force format even if a file system is found** all existing data on the volume will be overwritten. If the logical unit has existing data, and you do not use the **Force Format** option, the VMM job to assign the logical unit will complete with a warning. VMM assigns the logical unit to the host. You can format the disk later.

6. In **Mount Point**, select the mount options. Then click **OK** to assign the logical unit to the host.
7. VMM registers the storage logical unit to the host and mounts the storage disk.
 - To view the associated job information, open the **Jobs** workspace.
 - To verify that the logical unit was assigned, view the information on the **Storage** tab in the **Host Name > Properties** dialog box. The newly assigned logical unit appears under **Disk**. Click the new disk to view the disk details.
 - If the **Array** field is populated in the disk details, this indicates that the storage array is under VMM management.
8. To configure additional disk settings open Disk Management on the host. To open Disk Management, click **Start**, type **diskmgmt.msc** in the search box, and then press ENTER. The new disk appears in the list of disks as a basic disk. If you chose to format the disk, the disk is already formatted and online. You can right-click the disk to see the available options, such as **Format** and **Change Drive Letter and Paths**.

Configure storage for a Hyper-V cluster

1. Click **Fabric Servers > All Hosts**. Right-click the cluster you want to configure > **Properties**. In **Host Cluster Name > Properties** click a tab:
 - **Available Storage**: for adding available storage, converting available storage to shared storage (CSV), or removing available storage.
 - **Shared Volumes**: for adding cluster shared volumes (CSVs), converting CSVs to available storage, or removing CSVs. The cluster must run at least Windows Server 2012 to support CSVs.
2. Configure storage for the host cluster. Note that:
 - If you add available storage for CSVs, use only alphanumeric characters for a LUN. Note that you can't change the partition style of a disk that has already been initialized.
 - If you're converting available storage to CSVs, make sure there are no VMs on the cluster that have their associated .vhd or .vhdx files located on the storage that you want to convert.

Convert volumes one at a time. After conversion, confirm that the logical unit appears on the **Shared Volumes** tab.

Caution: If you convert shared to available storage, and the storage is being used by virtual machines, serious data loss can result.

- You can only remove storage if there are no VMs in the cluster currently using the storage for their vhds.
3. When you're ready to commit the changes, click **OK**.

Create an iSCSI session

1. On the target host, in the Services snap-in, make sure that the Microsoft iSCSI Initiator Service is started and set to Automatic.
2. In **Fabric > Servers > All Hosts Hosts**, right-click the host that you want to configure > **Properties**.
3. Under **iSCSI Arrays**, see if the storage array is already listed. If it is not, on the toolbar, next to **iSCSI Array**, click **Add**.
4. In the Create New iSCSI Session > **Array**, click the storage array you want to use.
5. Click **Create** to create a new session. Click Use advanced settings if you want to modify customized settings, including target listener, name, or the host NIC that you want to use.
6. The array that you added appears under **iSCSI Arrays**. Click the array to view more details.

Next steps

[Set up networking](#) for Hyper-V hosts and clusters.

Manage MPIO for Hyper-V hosts in the VMM fabric

12/7/2018 • 2 minutes to read

Read this article to learn about how System Center - Virtual Machine Manager (VMM) manages Multipath I/O (MPIO) on Hyper-V hosts.

When you add fiber channel or iSCSI storage to a Hyper-V host managed in the VMM fabric, the Multipath I/O (MPIO) feature must be enabled on each host.

- If MPIO is enabled on the host, VMM adds it for [supported storage arrays](#) using the Microsoft DSM. If you installed vendor-specific DSMs, then the vendor-specific MPIO settings will be used to connect to the storage array.
- If you add a host and MPIO isn't enabled, VMM will show a warning message in the **Jobs** window.
- If you add a host to VMM and enable MPIO afterwards, you need to add the MPIO feature, and then manually configure MPIO to add the discovered device hardware IDs. Alternatively, you can install vendor-specific DSMs.

MPIO tracking in VMM

When Hyper-V hosts and clusters are added to the VMM fabric, VMM deploys an agent to connectivity between the host and the VMM server. Additionally, VMM collects configuration information about the host or cluster and adds it to VMM. In the case of MPIO, VMM adds two registry keys containing MPIO information.

```
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\MPDEV\MPIOSupportedDeviceList
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\msdsm\Parameters\DsmSupportedDeviceList
```

After supported storage devices are added to the device list, VMM makes a "claim" on them, and a restart is required on the host. If you've added the host to VMM before deploying workloads on it, then this probably isn't an issue, but if workloads are already running on the host, this could cause interruptions. In an attempt to avoid potential outages, you can run a PowerShell script to prepopulate the MPIO registry keys on a host, before adding it to the VMM fabric. [Learn more](#) about this script.

Prevent VMM from claiming device IDs

If you don't want VMM to claim any storage device IDs for MPIO purposes, do the following:

1. Open registry location **HKLM\Software\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings**
2. In this location, create a registry key **RemoveMPIOHardwareIds** with type multi-string.
3. Add the device IDs from the default list. Make sure you use the same spacing.
4. Restart the VMM service.
5. Add the Hyper-V host in the VMM.

Next steps

[Provision a VM](#)

Manage port ACLs in VMM

5/30/2019 • 8 minutes to read

In System Center - Virtual Machine Manager (VMM), you can centrally configure and manage Hyper-V port access control lists (ACLs). Port ACLs can be configured for both a Network Controller- managed fabric, and for networks that aren't managed by Network Controller.

- A port ACL is a set of rules that filter traffic at the layer 2 port level. A port ACL in VMM filters access to a particular VMM object. A network object can have no more than one port ACL attached.
- An ACL contains rules, and can be attached to any number of network objects. You can create an ACL without rules, and then add rules at a later time. Each ACL rule corresponds to a only one port ACL
- If an ACL has multiple rules, they are applied based on priority. After a rule matches criteria and is applied, no other rules are processed.
- A global settings port ACL is applied to all VM virtual network adapters in an infrastructure. There's no separate object type for global settings. Instead, the global settings port ACL is attached to the VMM management server.
- Port ACL settings are exposed only through PowerShell cmdlets in VMM, and can't be configured in the VMM console.
- Port ACLs can be applied to:
 - Virtual subnets and adapters in a Network Controller deployment.
 - Virtual subnets, network adapters, VM networks, and the VMM management server in networks that aren't managed by Network Controller.

Before you start

- To apply an ACL to objects managed by Network Controller, you use the **ManagedByNC** flag, and set it to **True**. If it's not set to **True**, the ACL only applies to network objects that aren't managed by Network Controller.
- ACL types aren't interchangeable. You can't apply an ACL with **ManagedByNC** set to **false**, to objects managed by Network Controller, and vice versa.
- The key difference between these two kinds of ACLs is that you need to remediate each network adapter, after applying ACL on objects that aren't managed by Network Controller.
- There's also a difference in priority ranges:
 - **Hyper-V port ACLs (not managed by Network Controller):** 1 - 65535
 - **SDN port ACLs (managed by Network Controller):** 1 - 64500
- You need full VMM admin permissions to attach a port ACL to global settings. To attach the ACL to VMM objects (VM networks, subnets, virtual network adapters), you need to be a VMM admin or tenant admin, or a self-service user.

Unsupported scenarios

Here is a list of unsupported scenarios:

- Manage individual rules for a single instance, when the ACL is shared with multiple instances. All rules are managed centrally within their parent ACLs, and apply wherever the ACL is attached.
- Attach more than one ACL to an entity.
- Apply port ACLs to virtual network adapters in the Hyper-V parent partition (management operating system).
- Create port ACL rules in VMM that include IP-level protocols (other than TCP or UDP). Other protocols are still supported natively by Hyper-V.

- Apply port ACLs to logical networks, network sites (logical network definitions), subnet VLANs, and other VMM networking objects that aren't mentioned as specifically supported.

Deployment steps

Use the VMM PowerShell interface to do the following:

1. Define port ACLs and rules.
 - The rules are applied to virtual switch ports on Hyper-V servers as "extended port ACLs" (VMNetworkAdapterExtendedAcl). This means that they can apply only to hosts running Windows Server 2012 R2 or later, because VMM doesn't create legacy Hyper-V port ACLs (VMNetworkAdapterAcl) for earlier versions.
 - All port ACL rules defined in VMM are stateful for TCP. You can't create stateless TCP ACL rules.
2. Attach a port ACL to global settings. This applies the ACL to all VM virtual network adapters.
3. Attach the port ACLs to VM networks, VM subnets, or VM virtual network adapters.
4. Manage port ACL rules.

Create port ACLs

1. Open PowerShell in VMM.
2. Create a port ACL with the [New-SCPortACL](#) cmdlet.

```
New-SCPortACL [-Name] <String> [-Description <String>] [-JobVariable <String>] [-ManagedByNC] [-OnBehalfOfUser <String>] [-OnBehalfOfUserRole <UserRole>] [-Owner <String>] [-ProTipID <Guid>] [-RunAsynchronously] [-UserRole <UserRole>] [-VMMServer <ServerConnection>] [<CommonParameters>]
```

Parameters

PARAMETER	DETAILS
Name; Description	Port ACL name and description
JobVariable	Stores job progress
ManagedByNC	Specifies whether objects are managed by Network Controller
OnBehalfOfUser/OnBehalfOfRole	Run job with user name or role
Owner	Specifies the owner of a VMM object in the form of a valid domain user account. Example: Contoso\PattiFuller or PattiFuller@Contoso
ProTipID	ID of ProTip that triggered action
RunAsynchronously	Indicates whether job runs asynchronously
UserRole	Specifies user role
VMMServer	Specifies VMM server
CommonParameters	Learn more

Examples

Create a port ACL for objects managed by Network Controller "DemoACLManagedByNC":

```
PS: C:\> New-SCPortACL -Name "DemoACLManagedByNC" -Description "PortACL Example Managed by NC" -ManagedByN
```

Create a port ACL for objects not managed Network Controller "DemPortACL":

```
PS: C:\> New-SCPortACL -Name "DemoPortACL" -Description "Port ACL Example Non Managed by NC"
```

Define port ACL rules for a port ACL

1. Open PowerShell in VMM.
2. Create one or more rules with the [New-SCPortACLRule](#) cmdlet.

```
New-SCPortACLrule -PortACL <PortACL> -Name <string> [-Description <string>] -Type <Inbound | Outbound>
-Action <Allow | Deny> -Priority <uint16> -Protocol <Tcp | Udp | Any> [-LocalAddressPrefix <string:
IPAddress | IPSubnet>] [-LocalPortRange <string:X|X-Y|Any>] [-RemoteAddressPrefix <string: IPAddress |
IPSubnet>] [-RemotePortRange <string:X|X-Y|Any>]
```

Parameters

PARAMETER	DETAILS
Name, Description	Rule name and description
Type	Specifies the traffic direction for which the ACL is applied (Inbound or Outbound)
Action	Specifies whether the ACL allows or blocks traffic (Allow or Deny)
LocalAddressPrefix	Specifies the source IP address or subnet that's used to identify traffic that should be filtered.
LocalPortRange	Specifies the source port range that's used to identify traffic.
RemoteAddressPrefix	Specifies the destination IP address or subnet that's used to identify traffic for filtering.
RemotePortRange	Specifies the destination port range that's used to identify traffic.
Protocol	Specifies the protocol for which the rule is applied.
Priority	Specify the priority of the rule in in port ACL. Rules are applied according to order. Set a priority between 1 and 65535, where the lowest number has highest priority. Port ACLs rules for objects managed by Network Controller should be set equal to or greater than 100. Network Controller doesn't support priority below 100.

Examples

Create a port ACL and store the object in \$portACL:

```
PS: C:\> $portACL = New-SCPortACL -Name "RDP ACL" -Description "Acl on RDP access"
```

Create a port ACL rule to allow RDP access from a remote subnet:

```
PS: C:\> New-SCPortACLRule -Name "AllowRDPAccess" -PortACL $portACL -Description "Allow RDP Rule from a subnet" -Action Allow -Type Inbound -Priority 110 -Protocol Tcp -LocalPortRange 3389 -RemoteAddressPrefix 10.184.20.0/24
```

Modify the priority of an ACL rule:

```
PS: C:\> $portACLRule = Get-SCPortACLRule -Name "AllowRDPAccess" ``  
PS: C:\> Set-SCPortACLRule -PortACLRule $portACLRule -Priority 220
```

The first command gets the port ACL rule "AllowRDPAccess". The second command changes the priority of the rule to 220.

Modify the port ACL rule for the destination address range and protocol for a rule:

```
PS: C:\> $portACLRule = Get-SCPortACLRule -Name "AllowRDPAccess" ``  
PS: C:\> Set-SCPortACLRule -PortACLRule $portACLRule -RemoteAddressPrefix 172.185.21.0/24 -Protocol Udp
```

The first command retrieves rule "AllowRDPAccess". The second changes the protocol to UDP, and sets the destination to subnet 172.185.21.0/24.

Attach and detach port ACLs

A port ACL can be attached to global settings, VM networks, VM subnets, and virtual network adapters. A port ACL attached to global settings applies by default to all VM virtual network adapters.

1. Open PowerShell in VMM.
2. Attach a port ACL using the [Set-SCVMMServer](#) cmdlet.

```
Set-SCVMMServer -VMMServer <VMMServer> [-PortACL <NetworkAccessControlList> | -RemovePortACL ]
```

Parameters

PARAMETER	DETAILS
VMM server	Name of the VMM server on which the port ACL is applied.
PortACL	Optionally attaches the specified port ACL to global settings.

Examples

Attach an ACL to global settings:

```
Set-SCVMMServer -VMMServer "VMM.Contoso.Local" -PortACL $acl``  
ExampleL: `` Set-SCVMMServer -VMMServer "VMM.Contoso.Local" -PortACL $acl
```

Detach an ACL from global settings:

```
Set-SCVMMServer -VMMServer "VMM.Contoso.Local" -RemovePortACL
```

Attach an ACL to a VM network during creation:

```
New-SCVMNetwork [-PortACL <NetworkAccessControlList>] [rest of the parameters]
```

Attach an ACL to an existing VM network:

```
Set-SCVMNetwork -PortACL $acl`
```

Attach an ACL to a VM subnet during creation:

```
New-SCVMSubnet [-PortACL <NetworkAccessControlList>] [rest of the parameters]
```

Attach an ACL to an existing VM subnet:

```
Set-SCVMSubnet [-PortACL <NetworkAccessControlList> | -RemovePortACL] [rest of the parameters]
```

Retrieve and view port ACLs and rules

1. Open PowerShell in VMM.
2. Run the [Get-SCPortACL](#) cmdlet, to retrieve and view a port ACL:

```
Get-SCPortACL [[-Name] <String> ] [-ID <Guid> ] [-OnBehalfOfUser <String> ] [-OnBehalfOfUserRole <UserRole> ] [-VMMServer <ServerConnection> ] [ <CommonParameters>]
```

3. Run the [Get-SCPortACLRule](#) to retrieve and view a rule:

```
Get-SCPortACLRule [-Name <String> ] [-ID <Guid> ] [-OnBehalfOfUser <String> ] [-OnBehalfOfUserRole <UserRole> ] [-PortACL <PortACL> ] [-VMMServer <ServerConnection> ] [ <CommonParameters>]
```

Parameters

PARAMETER	DETAILS
No parameters	Retrieves all ACLs
Name/ID	Retrieve by name or GUID
OnBehalfOfUser/OnBehalfOfUserRole	Run with user name or role
VMMServer	Retrieve ACLs on specific VMM server
CommonParameters	Learn more

Examples

Retrieve a specific ACL:

```
PS: C:> $portACL = Get-SCPortACL -Name "DemoPortACL"
```

Get rules for a specific ACL:

```
PS: C:> Get-SCPortACLRule -Name "AllowRDPAccess"
```

Get all rules for ACL:

```
PS: C:> Get-SCPortACLRule -PortACL $portACL
```

Modify port ACLs and rules

1. Open PowerShell in VMM.
2. Run the [Set-SCPortACL](#) cmdlet, to modify a port ACL:

```
Set-SCPortACL [-PortACL] <PortACL> [[-Description] <String>] [-JobVariable <String>] [-Name <String>] [-OnBehalfOfUser <String>] [-OnBehalfOfUserRole <UserRole>] [-PROTipID <Guid>] [-RunAsynchronously] [-VMMServer <ServerConnection>] [<CommonParameters>]
```

3. Run the [Remove-SCPortACL](#) to remove an ACL:

```
Remove-SCPortACL [-PortACL] <PortACL> [-Confirm] [-JobVariable <String>] [-OnBehalfOfUser <String>] [-OnBehalfOfUserRole <UserRole>] [-PROTipID <Guid>] [-RunAsynchronously] [-VMMServer <ServerConnection>] [-WhatIf] [<CommonParameters>]
```

Parameters

PARAMETER	DETAILS
Name/Description	Name and description of port ACL
JobVariable	Stores job progress
OnBehalfOfUser/OnBehalfOfUserRole	Run with user name or role.
ProTipID	ID of ProTip that triggered action
RunAsynchronously	Indicates whether job runs asynchronously.
Confirm	Prompts before running job
WhatIf	Shows what happens without running command

Examples

Set an ACL description:

```
PS: C:> $portACL = Get-SCPortACL -Name "DemoPortACL"
```

```
PS: C:> Set-SCPortACL -PortACL $portACL -Description "Port ACL Example Non Managed by Network Controller"
```

The first cmdlet retrieves the ACL, the second sets the description on the ACL.

Remove an ACL:

```
PS: C:> $portACL = Get-SCPortACL -Name "DemoPortACL"
```

```
PS: C:> Remove-SCPortACL -PortACL $portACL
```

The first cmdlet retrieves the ACL, the second removes it.

Manage Hyper-V clusters in the VMM fabric

12/7/2018 • 3 minutes to read

Use this article to manage Hyper-V host clusters in the System Center - Virtual Machine Manager (VMM) fabric. You can configure cluster properties, and manage cluster nodes.

Configure cluster properties

1. In **Fabric**, right-click the cluster > **Properties**.
2. Configure the settings summarized in the table.

TAB	SETTINGS
General	<p>View the name, host group and description. You can also configure the Cluster reserve (nodes) setting, and view the cluster reserve state.</p> <p>The Cluster reserve (nodes) setting specifies the number of node failures a cluster must be able to sustain while still supporting all virtual machines deployed on the host cluster. If the cluster cannot withstand the specified number of node failures and still keep all of the virtual machines running, the cluster is placed in an over-committed state. When over-committed, the clustered hosts receive a zero rating during virtual machine placement. An administrator can override the rating and place a highly-available virtual machine on an over-committed cluster during a manual placement.</p>
Status	<p>View detailed status information for the host cluster:</p> <p>Cluster validation test runs and successes. Includes a link to the latest validation report (if available). Note that accessing the report requires administrative permissions on the cluster node where the report is located. For host clusters, you can perform an on-demand cluster validation through VMM. To do this, in the Fabric workspace, locate and click the host cluster. Then, on the Host Cluster tab, click Validate Cluster. Cluster validation begins immediately.</p> <p>Online elements in the cluster: cluster core resources, disk witness in quorum, and the cluster service on each node.</p>
Available Storage	<p>Shows available storage, that is, storage logical units that are assigned to the host cluster but are not Cluster Shared Volumes (CSV).</p> <p>You can also do the following:</p> <p>Add and remove storage logical units that are managed by [VMM].</p> <p>Convert available storage to shared storage (CSV).</p>

TAB	SETTINGS
Shared Volumes	<p>Shows the shared volumes (CSVs) that are allocated to the host cluster. You can also do the following:</p> <p>Add and remove CSVs that are managed by VMM.</p> <p>Convert CSVs to available (non-CSV) storage.</p>
Custom Properties	Custom properties that you manage.

Add a node to the cluster

1. If you already used Failover Cluster Manager to add the node, then in **Fabric > Servers > All Hosts**, right-click the host with a **Pending** status, and click **Add to Host Cluster**.
2. If you didn't add the node with the Failover Cluster Manager, you can add hosts that are already managed by VMM. In **Fabric > Servers > All Hosts**, right-click the cluster > **Add Cluster Node**. In the Add Nodes Wizard > **Resource Type**, select the RunAs account that will be used to add the nodes. Make sure **Existing servers running a Windows Server operating system** is selected. In **Select Hosts** select the Hyper-V host server that you want to add. Finish the wizard and verify the settings.

Remove a node from the cluster

1. Click **Fabric > Servers > All Hosts**.
2. Locate the cluster node you want to remove and view the status in the **Hosts** pane.
3. If the node isn't in maintenance mode click Start Maintenance Mode. Click **Move all virtual machines to other hosts in the cluster** and verify the status.
4. Right-click the host > **Remove Cluster Node > Yes**. During the job to remove the node any shared storage is unregistered from the node. If you manage storage outside VMM then you should unregister the storage from the node.

Uncluster a cluster

Remove a host cluster as follows:

1. Click **Fabric > Servers > All Hosts**.
2. Make sure the cluster isn't supporting any highly-available VMs or clustered services/apps.
3. Right-click the host cluster > **Uncluster**. Click **Yes** to continue.
4. During the job to remove the cluster any shared storage is unregistered from the cluster nodes. If you manage storage outside VMM, then you should unregister the storage.

Update Hyper-V hosts and clusters

3/19/2019 • 2 minutes to read

Read this article to learn about keeping Hyper-V hosts and clusters updated in the System Center Virtual Machine Manager (VMM) fabric.

Before you start

You'll need to set up an update (WSUS server) in the VMM fabric, and configure update baselines.

Update a host or cluster

1. In **Fabric > Servers > Show > Compliance**, select the computers you want to update. All the baselines for the computer will be shown. The computer could be compliant for some baselines and not for others.
2. Click **Remediate**. You'll only see this option if the selected objects aren't compliant.
3. In **Update Remediation**, select or clear update baselines or individual updates to determine which updates to install. When you select a computer all updates are initially selected. If you're running a Hyper-V cluster note that:
 - To update all cluster nodes, select the host cluster by its cluster name.
 - To update single nodes, select the individual hosts in the cluster. With this setting VMM doesn't display cluster remediation options but treats each node as a standalone Hyper-V host.
4. If you prefer to restart the computers manually after remediation completes, click **Do not restart the servers after remediation**. By default, the computer restarts if any updates require it. If you choose not to restart and updates need it, the computer status will be **Pending Machine Reboot** after the remediation. The updates won't be activated until you restart. With this status VMM won't scan the machines for compliance during refreshes.
5. If you're running a cluster:
 - If cluster node is already in maintenance mode, select **Allow remediation of clusters with nodes already in maintenance mode**. In maintenance mode VMs can't be created or moved to the host. The host has a zero rating and is excluded from dynamic optimization.
 - Select **Live migration** to remove virtual machines from a host before performing update remediation so that they remain online. If you don't need to keep them online and want to perform a quicker update, click **Save state** to shut down virtual machines and proceed with remediation.
6. Click **Remediate** to start updating. After remediation, if no reboot is pending the server or cluster will show as **Compliant**.

NOTE

Remediation of single node in a S2D cluster might not succeed if remediation is done by selecting the host. We recommend you to remediate at cluster level and then select the hosts for remediation.

In S2D cluster, once the update is installed on the node, the node is rebooted and repair (resync) operation is initiated. If the resync is not completed within 600 minutes (default value), the update job fails. To avoid the job failure, you can change the wait time in the registry key.

- Minimum value for the resync wait time is 600 minutes and the maximum value can be of 2880 minutes.
- **S2DResyncWaitTimeInMins** is the registry key setting that you can use to change the default Resync wait

time.

- *Here is the path for the registry key:*

\HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine Manager
Server\Settings > RegistryKey

Next steps

[Learn about](#) VMM updates

Perform a rolling upgrade of a Hyper-V host cluster to Windows Server in VMM

8/29/2019 • 3 minutes to read

Cluster rolling upgrade was introduced in Windows Server 2016. This feature enables you to upgrade the operating system of cluster nodes without stopping Hyper-V workloads running on the nodes. [Read more](#) about rolling upgrade requirements and architecture.

NOTE

System Center 2019 - Virtual Machine Manager (VMM) supports rolling upgrade of a Hyper-V host cluster from Windows Server 2016 to Windows Server 2019. Versions earlier to VMM 2019 supports rolling upgrade to 2016 from 2012 R2. Use the following procedures as applicable for the version of VMM you are using and the upgrade version it supports.

Cluster rolling upgrade in VMM

System Center - Virtual Machine Manager (VMM) supports using the rolling upgrade feature to upgrade Hyper-V clusters in the VMM fabric. You can upgrade an entire cluster or specific cluster nodes. Here's what the upgrade does:

- **Creates a template:** Creates a template of the node configuration by combining the appropriate physical computer profile with the node configuration settings detailed in the upgrade wizard.
- **Migrates workloads:** Migrates workloads off the node so workload operations aren't interrupted.
- **Removes node:** Puts the node into maintenance mode and then removes it from the cluster. This removes all VMM agents, virtual switch extensions, and so forth from the node.
- **Provisions the node:** Provisions the node running Windows Server 2016/2019, and configures it according to the saved template.
- **Returns the node to VMM:** Brings the node back under VMM management and installs the VMM agent.
- **Returns the node to the cluster:** Adds the node back into the cluster, brings it out of maintenance mode, and returns virtual machine workloads to it.

NOTE

Ensure to install the latest updates on the VHD that you want to use as a physical computer profile.

Before you start

Review the platform [restrictions and limitations](#) before you start Cluster rolling upgrade.

- The cluster must be managed by VMM.
- The cluster must be running Windows Server 2012 R2 (in case of releases prior to 2019) or Windows Server 2016 (in case of VMM 2019).
- The cluster must meet the [requirements](#) for bare metal deployment. The only exception is that the physical computer profile doesn't need to include network or disk configuration details. During the upgrade VMM records the node's network and disk configuration and uses that information instead of the computer profile.
- You can upgrade nodes that weren't originally provisioned using bare metal as long as those nodes meet bare metal requirements such as BMC. You'll need to provide this information in the upgrade wizard.

- The VMM library needs a virtual hard disk configured with Windows Server 2016.

Run the upgrade

1. Click **Fabric > Servers > All Hosts**. Right-click the host cluster > **Upgrade Cluster**.
2. In the Upgrade Wizard > **Nodes**, click the nodes you want to upgrade or **Select All**. Then click **Physical computer profile**, and select the profile for the nodes.
3. In **BMC Configuration**, select the Run As account with permissions to access the BMC or create a new one. In **Out-of-band management protocol** click the protocol that the BMCs use. To use DCMI click IPMI. DCMI is supported even though it's not listed. Make sure the correct port is listed.
4. In **Deployment Customization**, review the nodes to upgrade. If the wizard couldn't figure out all of the settings it displays a **Missing Settings** alert for the node. For example if the node wasn't provisioned by bare metal BMC settings might not be complete. Fill in the missing information.
 - Enter the BMC IP address if required. You can also change the node name. Don't clear **Skip Active Directory check for this computer name** unless you're changing the node name and you want to make sure the new name is not in use.
 - In the network adapter configuration you can specify the MAC address. Do this if you're configuring the management adapter for the cluster, and you want to configure it as a virtual network adapter. It's not the MAC address of the BMC. If you choose to specify static IP settings for the adapter, select a logical network and an IP subnet if applicable. If the subnet contains an address pool you can select **Obtain an IP address corresponding to the selected subnet**. Otherwise type an IP address within the logical network.
5. In **Summary** click **Finish** to begin the upgrade. If the wizard finishes the nodes upgrade successfully, the wizard upgrade the cluster functional level to Windows Server 2016/2019.

If for some reason you need to update the cluster functional level of a cluster that was upgraded outside VMM, right-click the **Cluster > Update Version**. This could happen if you upgrade the cluster nodes before adding the cluster to the VMM fabric, but the cluster is still functioning as a Windows Server 2012 R2/2016 cluster.

Service hosts and virtual machines in the VMM compute fabric

12/7/2018 • 2 minutes to read

Read this article to learn about service Hyper-V hosts and virtual machines in the System Center Virtual Machine Manager (VMM) fabric.

Before you start

We recommend you keep hosts in a Hyper-V host cluster up-to-date with the same updates, to avoid VMM operations issues.

You can service hosts and VMs by:

- **Setting up servicing windows:** Servicing windows provide a method for scheduling servicing outside VMM. You can associate a servicing window with individual hosts, virtual machines, or services. Before using other applications to schedule maintenance tasks, you can use Windows PowerShell scripts or custom applications to query the object and determines if it is currently in a servicing window. Servicing windows do not interfere with the regular use and functionality of VMM.
- **Placing hosts in maintenance mode:** You can start maintenance mode for a virtual machine host whenever you need to perform maintenance tasks on the physical host, such as applying security updates or replacing hardware on the physical host computer. You can place Hyper-V hosts and VMware ESX hosts in the VMM fabric into maintenance mode. Note that:
 - When a host is in maintenance mode, the following restrictions apply:
 - VMs can't be created on the host.
 - VMs can't be moved to the host.
 - The host has a zero rating and can't be selected for placement.
 - The host is excluded from dynamic optimization.
 - For ESX hosts, if the VMware Distributed Resources Scheduler is not configured, all virtual machines on the host must be either manually shut down or moved to another host to successfully start maintenance mode on an ESX host.
 - If you're monitoring with Operations Manager and you have maintenance mode enabled for monitoring, the Operations Manager agent suppresses alerts, notifications, rules, monitors, automatic responses, state changes, and new alerts for hosts that are placed in maintenance mode.

Set up a servicing window

1. In the VMM console, click **Settings > Create > Create Servicing Window**.
2. In **New Servicing Window**, specify a name and optional description for the window.
3. In **Category**, enter or select the category of servicing window.
4. In **Start time**, enter the date, time of day, and time zone for the maintenance window.
5. In **Duration**, specify the number of hours or minutes in the servicing window.
6. Under **Recurrence pattern**, select the frequency (Daily, Weekly, or Monthly), and then schedule the occurrences within that frequency.
7. After the window is set up, you can assign it to a host or VM. To assign to a host, click the host properties > **Servicing Window > Manage**, and select the window you want to add to the host.

Put hosts in maintenance mode

1. In the VMM console, click **Fabric > Fabric Resources > Servers > All Hosts**.
2. Select the host to place in maintenance mode, and in the **Host** group click **Start Maintenance Mode**. You can select to:
 - You can select **Move all virtual machines to other hosts in the cluster** if you want to move all highly available VMs to other hosts in the cluster (the host must be in a cluster that's capable of live migration)
 - Otherwise, select **Place all running virtual machines into a saved state**. Note that list causes a loss of service for users currently using the VM.
3. You can verify that host is in maintenance mode by checking its status in **Fabric > Hosts**.

To bring a host out of maintenance node, select it and click **Stop Maintenance Mode**. Note that VMM doesn't automatically restart the VM, and doesn't automatically migrate VMs back to the host.

Set up VMware servers in the VMM compute fabric

12/7/2018 • 9 minutes to read

Read this article to learn about managing VMware servers and VMs in the System Center - Virtual Machine Manager (VMM) fabric.

VMM enables you to deploy and manage virtual machines and services across multiple hypervisor platforms, including VMware vSphere hosts and vCenter servers.

- You can add vCenter and vSphere hosts to the VMM fabric.
- VMM integrates directly with VMware vCenter Server. Through the VMM console, you can manage the day-to-day operations of VMware vSphere hosts and clusters, such as the discovery and management of hosts.
- VMM provides the ability to create, manage, store, place and, deploy virtual machines on vSphere hosts. You can import VMware templates.
- You can associate host adapters with VMM logical networks. More advanced management takes place on the vCenter Server, for example configuring port groups, standard and distributed virtual switches (or "vSwitches"), vMotion and, Storage vMotion.
- You can convert VMware VMs to Hyper-V.

Before you start

- VMM supports the management of hosts and clusters running VMware. [Learn more](#) for supported versions of VMware.
- You need a vCenter server in your deployment. vSphere hosts and host clusters should be managed by a vCenter server, which in turn is managed by VMM.
- The following features are supported by VMM when hosts and clusters are managed with a vCenter server:
 - VMM command shell (same shell across all hypervisors)
 - VM placement based on host ratings when you create, deploy, or migrate VMware VMs. Includes concurrent VM deployment during service deployment.
 - You can deploy VMM services to vSphere hosts. You can't deploy vApps.
 - You can make vSphere host resources available to a VMM cloud by creating clouds from host groups in which vSphere hosts are located, or by creating a cloud from a VMware resource pool. Note that VMM doesn't integrate with VMware vCloud.
 - You can use dynamic optimization and power optimization for vSphere hosts. VMM can load balance virtual machines on vSphere clusters using live migration. With power optimization, you can configure VMM to turn vSphere hosts on and off for power management.
 - You can transfer VMware resources using live migrating between hosts in a cluster (uses vMotion) and live storage migration (uses storage vMotion). Resources supported for transfer include network migration to and from the library and between hosts. Note that VMware thin provision disks become thick when you migrate a disk to the VMM library.
 - You can place vSphere hosts managed by VMM into and out of maintenance mode.
 - You can organization and store VMware VMs, VMDK files, and VMware templates in the VMM library. You can create new VMs from templates. Note that VMM doesn't support older VMDK disk types. These disk types are supported: regular VMDK files (VMFS and monolithic flat), VMDK files that are used to access physical disks (vmfsPassthroughRawDeviceMap), and snapshots (vmfssparse).
 - You can create templates using .vmdk files stored in the library. You can also import templates stored on vSphere hosts (only template metadata is imported to VMM)
 - VMM supports existing standard and distributed vSwitches and port groups. vSwitches and port

groups must be configured with the vCenter server.

- You can do regular VMM networking tasks, including assigning logical networks, static IP address and MAC address to Windows-based VMs running on VMware.
- VMM supports and recognized VMware Paravirtual SCSI (PVSCSI) storage adapters.
- VMM doesn't support VMware VM with virtual hard disks connected to an IDE bus.
- VMM supports VMware thin provision hard disk through the dynamic disk type. Note that if you create and deploy a VM to a vSphere host configured to use a dynamic disk the disk will be thin provisioned. If a VM was created as a thin provisioned disk out-of-band VM will display it as dynamic. If you save a thin provision disk to the library VMM will save it as thick. It remains thick if you create a VM from it.
- All storage must be added to vSphere hosts outside VMM.
- Communication between VMM and the vCenter server is SSL encrypted. You'll need a certificate to identify the vCenter server You can use a self-signed certificate for a vCenter server or a third-party verified certificate.
- If you're using a self-signed certificate for authenticating the vCenter server to VMM you can manually import the certificate to the Trusted People certificate store on the VMM management server before adding the vCenter server. If you don't you'll be prompted to do this during deployment.
- You'll need an account with admin permissions for the vCenter server (local or Active Directory account) and an account with admin permissions for the vSphere hosts. You can create Run As accounts before you begin. If you don't you'll need to create accounts during the deployment procedure.
- You can decide whether you want VMM to communicate with the vSphere hosts managed by the vCenter server over a secure connection. If so you'll need a certificate to authentication communications on each vSphere host or cluster. You can either use the self-signed certificate that VMware created when vSphere was installed on the host, or a certificate from a trusted CA. If you're using a self-signed certificate you can import it from each vSphere host to the vMM management server before you begin deployment
- Before you configure network settings for vSphere hosts make sure you've created logical networks that you want to associate with the physical network adapters on the hosts.

Add a v-Center server

1. Click **Fabric > Servers > vCenter servers > Add > Add Resources > VMware vCenter Server**.
2. In **Add VMware vCenter Server** specify the name (FQDN, NetBIOS, or IP address) of the vCenter server. Add the port needed to connect to the vCenter server (443 by default).
3. In **Run As account** click the Run As account with admin permissions for the vCenter server. Click **Create Run As Account** if you don't have one.
4. In **Security** select or clear **Communicate with VMware ESX hosts in secure mode**. We recommend you keep the setting checked. If selected you'll need a certificate and public key for each vSphere host managed by the vCenter server.
5. If you're using a self-signed certificate to communicate with the vCenter server and you haven't manually copied it to the Trusted People certificate store the **Import Certificate** dialog will appear. Click **Import** to add the certificate to the store.
6. In **Jobs**, wait until the job has a Completed status and then check that the server appears in **Fabric > Servers > vCenter Server** with a **Responding** status.

Add an ESX/ESXi host

1. Ensure that the vCenter server is managed by VMM before you start. When you add the vCenter server, vSphere hosts for the server are discovered automatically.
2. Click **Fabric > Add Resources > VMware ESX Hosts and Clusters**.
3. In the **Add Resource Wizard > Credentials**, click the Run As account that has admin permissions on the

vSphere hosts you want to add. Create a Run As account if you don't have one

4. In **Target Resources** click the vCenter server. If the hosts are clustered they'll be listed together with the cluster nodes.
5. In **Computer Name**, select the hosts or clusters you want to add, or **Select All**.
6. In **Host Settings**, click the host group to which you want to assign the host or cluster. You don't need to add VM placement paths.
7. In **Summary**, verify the settings and click **Finish**. Wait until the Jobs dialog shows a **Completed** status.
8. Click **Fabric > Servers > All Hosts** and in the host group check the status of each host or cluster. Either **OK** or **OK (limited)**.
9. If the status is limited it means you've enabled the setting **Communicate with VMware ESX hosts in secure mode** but haven't yet imported a certificate from each vSphere host. To modify the security setting right-click the vCenter server > **Properties > Security**.
10. To import the certificate click each relevant host name > **Properties > Management > Retrieve > OK**. The host status should be **OK** after the import.

Associate host adapters with logical networks

By default when you added vSphere hosts to VMM, VMM automatically created logical networks that match the virtual network switch name. Note that VMM doesn't automatically create port groups so you'll need to configure port groups with the necessary VLANs that correspond to network sites on the vCenter server.

Associate the logical network with the physical network adapter (for an external virtual network) as follows:

1. Click **Fabric > Servers > All Hosts > vSphere host > Host > Properties > Hardware**.
2. In **Network Adapters**, select the physical network adapter on the host. In **Logical network connectivity**, select the logical networks you want to associate with the adapter. Note that only logical networks available to the host group are available.
3. Click **Advanced > Advanced Network Adapter Properties** to see IP subnets and VLANs available for a logical network. By default for a logical network the subnets and VLANs are scope to the host group or inherited via a parent host group. If none appears it indicates that no network site existing for the logical network. If **Unassigned** is available click it to view VLANs to which the physical adapter is connected, but that aren't included in a network site.
4. View virtual network settings in the host properties > **Virtual Networks**. View compliance information in **Fabric > Networking > Logical Networks > Hosts > Logical Network Information for Hosts > Compliance**. A status of **Fully compliant** indicates that all subnets and VLAN that are in the network site are assigned to the network adapter.

Import templates from vCenter

You can import VMware templates from the vCenter server to the VMM library. VMM copies only the metadata associated with the template and not the .vmdk file. This means that VMM is dependent on the vCenter server to use the template.

1. Click **Library > Home > Import > Import VMware template**.
2. In **Import VMware Templates**, select each template you want to import and click OK.
3. Verify the templates in **Library > Templates > VM Templates**.

Set up a servicing window for a VMware host

Servicing windows provide a method for scheduling servicing outside VMM. You can associate a servicing window with individual hosts, virtual machines, or services. Before using other applications to schedule maintenance tasks, you can use Windows PowerShell scripts or custom applications to query the object and

determines if it is currently in a servicing window. Servicing windows do not interfere with the regular use and functionality of VMM. Set up a servicing window as follows:

1. In the VMM console, click **Settings > Create > Create Servicing Window**.
2. In New Servicing Window, specify a name and optional description for the window.
3. In **Category**, enter or select the category of servicing window.
4. In **Start time**, enter the date, time of day, and time zone for the maintenance window.
5. In **Duration**, specify the number of hours or minutes in the servicing window.
6. Under **Recurrence pattern**, select the frequency (Daily, Weekly, or Monthly), and then schedule the occurrences within that frequency.
7. After the window is set up, you can assign it to a host or VM. To assign to a host, click the host properties > **Servicing Window > Manage**, and select the window you want to add to the host.

Set up infrastructure servers in the VMM compute fabric

12/7/2018 • 2 minutes to read

Read this article to learn about adding and managing infrastructure servers in the System Center - Virtual Machine Manager (VMM) fabric.

In addition to the infrastructure servers used by the VMM fabric ([library server](#), [PXE servers](#), [IPAM servers](#)), you can add other infrastructure servers such as Active Directory, DNS, DHCP, System Center to the VMM fabric. This allows you to manage and update all of these servers in the same location.

The **Infrastructure** node in the VMM console shows infrastructure servers you add, as well as the VMM management server, vCenter servers, VMM library servers, IPAM and PXE servers if you add them.

Add an infrastructure server

1. Click **Fabric > Servers > Infrastructure > Add Resources > Infrastructure Server**.
2. In **Add Infrastructure Server Wizard** specify the FQDN of the server you want to add an account with permissions for the server. Then click **Add**.

Update infrastructure servers

In order to update infrastructure servers you'll need to [set up a WSUS server and configure update baselines](#).

After the WSUS server is in place you can update infrastructure servers as follows:

1. In **Fabric > Servers > Home > Show > Compliance** select the server you want to update. All the baselines for the server will be shown. The server could be compliant for some baselines and not for others.
2. Click **Remediate**. You'll only see this option if the selected objects aren't compliant.
3. In **Update Remediation** select or clear update baselines or individual updates to determine which updates to install. When you select a computer all updates are initially selected.
4. If you prefer to restart the server manually after remediation completes click **Do not restart the servers after remediation**. By default, the server restarts if any updates require it. If you choose not to restart and updates need it the computer status will be **Pending Machine Reboot** after the remediation. The updates won't be activated until you restart. With this status VMM won't scan the machines for compliance during refreshes.
5. Click **Remediate** to start updating.

Set up update servers in the VMM compute fabric

2/25/2019 • 8 minutes to read

Learn about deploying update servers in the System Center - Virtual Machine Manager (VMM) fabric.

You deploy update servers in the VMM fabric to manage compliance and remediation for virtualization hosts, library servers, the VMM management server, PXE servers, the WSUS server itself, and any infrastructure servers.

[Learn more.](#)

This article explains the prerequisites, information about how to add a WSUS server to the fabric, set up update baselines, run a scan, and create update exemptions.

Before you start

- The WSUS server must be running Windows Server Update Service (WSUS) 4.0 or later and supported versions of Windows Server, as detailed in [system requirements](#).
- The WSUS server must be in the same domain as the VMM server, or in a domain with full trust.
- VMM can use a WSUS root server or downstream WSUS server. You can't use a WSUS replica server.
- The WSUS server can be dedicated to VMM or an existing server.
- The WSUS server can be installed on the VMM management server but if you'll be processing a large number of updates we recommend a separate server.
- VMM can work with System Center Updates Publisher but only full content updates are supported. Metadata-only updates can't be added to a baseline.
- After you add a WSUS server to VMM you should manage it in the VMM console and not in the WSUS console. In VMM, you update the properties of the update server to configure a proxy server for synchronizations and to change the update categories, products, and supported languages that are synchronized by the WSUS server.
- In VMM, administrators and delegated administrators manage fabric updates. Only administrators can manage the update server and synchronize updates. Delegated administrators can scan and remediate updates on computers that are within the scope of their user roles. Delegated administrators can use baselines created by administrators and other delegated administrators. But delegated administrators cannot modify or delete baselines created by others.

Add a WSUS server to the VMM fabric

1. Ensure that the server is running the WSUS role.
2. Click **Fabric > Home > Add > Add Resources > Update Server**.
3. In **Add Windows Server Update Services Server**, specify the name, port, and credentials of the WSUS server. The account needs admin rights on the server. Use an existing Run As account or create a new one. Specify whether you want to use SSL for connections.
4. The WSUS server will be added to the fabric, followed by initial synchronization of the updates catalog. This could take a while. Monitor status in the **Add Update Server** and **Synchronize Update Server** jobs.
5. After the server is added, you can update its properties to configure a proxy server for synchronization. In **Fabric > Servers > Update Server > Properties**, click the **Proxy Server** tab. Configure WSUS to use a proxy server when configuring updates, or update the port for an existing proxy server.
6. In addition, you can click **Update Classification** to select the update classification that you want to synchronize, click **Product** to select the products that want to include when synchronizing, and **Language** to select the supported synchronization languages.

After you've added the server you can update the WSUS settings and perform manual synchronization in **Servers** > WSUS server name > **Update Server**.

Add WSUS servers that are managed in Configuration Manager

If you want to add an existing WSUS server from a Configuration Manager environment to the VMM fabric you'll need to do the following:

1. Create a collection in Configuration Manager and add any servers that you'd like to add to the VMM fabric.
2. Exclude this collection from any software update deployments delivered by Configuration Manager. This ensures that VMM controls update management for the servers. You'll still be able to view compliance information for the collection in Configuration Manager reports.
3. If you want to include VMM compliance information in Configuration Manager, create an update group in Configuration Manager that contains all of the updates for which you want to measure compliance for the machines that will be in the VMM fabric. This update group is only for reporting. Don't deploy it to the machines managed by VMM.
4. Now add the WSUS server as described above.
5. After adding the server, select **Update Server** > **Properties** > **General** > **Allow Update Server configuration changes**.

Create and assign update baselines

After you've added the WSUS server to the fabric you can configure update baselines. An update baseline contains a set of required updates scoped to an object such as a host group, a standalone host, a host cluster, a VMM management server, or an infrastructure server.

- Update baselines can be assigned to host groups and to individual computers based on their role in VMM.
- Update baselines that are assigned to a host group are applied to all stand-alone hosts and host clusters in the host group, as well as the stand-alone hosts and host clusters in child host groups.
- During a compliance scan, computers that are assigned to a baseline are graded for compliance with their assigned baselines. After a computer is found noncompliant, an administrator brings the computer into compliance through update remediation.
- If a host is moved from one host group to another, the baselines for the new host group are applied to the host, and the baselines for the preceding host group no longer apply - that is, unless the baseline is assigned to both host groups. Explicit baseline assignments to a managed host stay with the host when it is moved from one host group to another. It is only when the baseline is assigned to a host group that baseline assignments get revoked during the move.
- You can use two methods to prepare update baselines for remediation:
 - A VMM built-in update baseline: Sample Baseline for Critical Updates and Sample Baseline for Security Updates.
 - A custom update baseline.

Assign servers to a built-in baseline

1. Click **Library** > **Update Catalog and Baselines** > **Update Baselines**.
2. In **Baselines**, click the baseline you want to use.
3. Click **Home** > **Properties** > **Updates** for the baseline. In **Updates**, add or remove baselines as required. To ensure all security updates are remediated don't remove anything.
4. Click **Assignment Scope**, and select the host groups, clusters, standalone servers, and infrastructure servers to add to the baseline. Or click **All Hosts** to add all.

NOTE

We recommend that you do not update the servers in a Storage Spaces Direct cluster (either hyper-converged or disaggregated) by using VMM, as it might cause data loss.

Assign servers to a custom baseline

1. Click **Library > Update Catalog and Baselines > Update Baselines**.
2. Click **Home > Create > Baseline** for the baseline.
3. In **Update Baseline Wizard > General**, specify a name and description.
4. In **Update**, add the updates you want to include.
5. In **Assignment Scope**, expand **Host Groups** and **Infrastructure**. Select the groups and servers you want to add.
6. In **Summary**, click **Finish**, and accept the **Microsoft License Terms** if needed to install any of the updates. Verify the baseline in **Library > Update Catalog and Baselines > Baselines**.

NOTE

We recommend that you do not update the servers in a Storage Spaces Direct cluster (either hyper-converged or disaggregated) by using VMM, as it might cause data loss.

Scan for update compliance

After you assign computers to an update baseline you can scan them to determine their compliance status for the baselines. When a computer is scanned for compliance.

- WSUS checks each update in the assigned update baselines to determine whether the update is applicable and if the update is applicable, whether the update has been installed.
- After a compliance scan, for every computer, each update has a compliance status of Compliant, Non Compliant, Error, Pending Reboot, or Unknown. You can view compliance properties for additional information.
- The compliance scan focuses only on the updates that the administrator has identified as important by adding them to a baseline. That enables organizations to monitor for compliance for what is deemed important for their organization.
- The following changes can cause an Unknown update status for a computer, and should be followed by a scan operation to access the computer's compliance status:
 - A host is moved from one host group to another host group.
 - An update is added to or removed from a baseline that is assigned to a computer.
 - The computer is added to the scope of a baseline.

To check compliance:

1. Click **Fabric > Servers**
2. In **Home > Show** click **Compliance**.
3. Since you haven't yet scanned computers the compliance status will show as **Unknown**, with an operational status **Pending Compliance Scan**.
4. Select the computers you want to check, and click **Scan**.
5. While the scan is in progress status will be Unknown. After it's finished compliance status for each update will be **Compliant**, **Non-Compliant**, or **Error**.

Manage update exemptions

You can create update exemptions for specific machines. For example if an update has caused the machine to be in

an unhealthy state, you could uninstall the update out-of-band and then exempt the machine from the update until the issue is resolved. When the compliance scan next runs the machine will show as **Non Compliant**.

1. Click **Fabric > Home > Show > Compliance**. Then on the **Fabric** node, click **Servers** and navigate to the server you want to exempt.
2. In the result pane expand the update baselines for the machine and click the update to select it.
3. Click **Compliance > Compliance Properties**.
4. In **Compliance Properties**, select the update > **Create**.
5. In **Create Exemption**, add notes about the reason and the expected exemption data. Change the update status to **Exempt**.
6. After you've resolved the issue and you want to cancel the exemption so that the machine is compliant again, in **Compliance Properties** select the exemption > **Delete > Yes**.
7. To return the server to a compliant state, select the non-compliant server and click **Remediate** on the **Compliance** tab.

Set up the VMM networking fabric

12/7/2018 • 2 minutes to read

This article provides an overview of setting up the System Center - Virtual Machine Manager (VMM) networking fabric.

Here's what you'll typically do to set up networking in the VMM fabric:

1. [Set up logical networks](#): Create logical networks that maps to your physical networks. You can create network sites that map to network sites in your physical network. For example IP subnets, VLNS, or subnet/VLAN pairs. Then if you're not using DHCP you create IP address pools for the network sites that exist within your physical networks.
2. [Create VM networks](#): Create VM networks that map to network sites that exist within your physical networks.
3. [Set up IP address pools](#): Create address pool to allocate static IP addresses. You'll need to configure pools for logical networks, and in some circumstances for VM networks too.
4. [Add a gateway](#): You might need to set up network virtualization gateways in the VMM networking fabric. By default, if you're using isolated VM networks in your VMM fabric, VMs associated with a network can only connect to machines in the same subnet. If you want to connect VMs further than the subnet you'll need a gateway.
5. [Create port profiles](#): Create uplink port profiles that indicates to VMM which networks a host can connect to on a specific network adapter. If required create virtual port profiles to specify settings that should be applied to virtual network adapters. You can create custom port classifications to abstract virtual port profiles.
6. [Configure logical switches](#): Create a logical switch, apply it to a host and select the network adapters on the host that you want to bind to the switch. When you apply the switch networking settings will be applied to the host.

Set up logical networks in the VMM fabric

12/7/2018 • 6 minutes to read

Read this article to learn how to create logical networks in System Center - Virtual Machine Manager (VMM).

You have different types of networks in your organization, including corporate networks, management networks, and others. In VMM each of these networks is defined as a logical network. Logical networks are logical objects that mirror your physical networks.

When you create logical networks, you assign them properties that match your physical environment. You specify the type of network, the associated network sites associated, and the static address pools if you're not using DHCP to assign IP addresses to VMs you create in the network sites. You also specify whether networks are isolated physically or virtually, using network virtualization and virtual LANs (VLANs).

You use logical networks when you provision virtualization hosts in the VMM fabric. You associate physical adapters on hosts with logical networks.

VMM virtual machine (VM) networks are based on logical networks. VM networks provide an interface through which VMs connect to a logical network. A logical network can have a single, or multiple, VM networks mapped to it.

Before you start

Before you start, it's important to understand how logical networks work in VMM.

- **Automatic logical networks:** By default, VMM creates logical networks automatically. When you provision a host in the VMM fabric and there's no VMM logical network associated with a physical network adapter on that host, VMM automatically creates a logical network and associates it with an adapter. By default for the logical network VMM first DNS suffix label of the connection-specific DNS suffix. By default VMM also creates a VM network configured with **No isolation**.
- **Manual logical networks:** When you create a logical network manually you specify:
 - **Network type:** You specify whether the network is isolated or not, and if it is how it's isolated. Then when you create VM networks based on the logical network they'll be created with the type of network you specified.
 - **No isolation:** This is the simplest type of network model that specifies there's just a single network within which machines can connect to each other with no need to isolate these machines from each other. VM networks in VMM provide an interface through which VMs connect to a logical network, and in a no isolation model you'll have a single VM network mapped to a logical network.
 - **Isolation:** More often you'll want to isolate networks from each other. For example you might want to isolate networks that have different purposes, or you might be a provider who wants to host workloads for multiple tenants on a single logical network, with isolation for each tenant. In this case you'll have multiple VM networks mapped to a logical network. VM networks mapped to a logical network can be isolated using VLANs/private VLANs, or network virtualization. Note that:
 - A typical setup might be an infrastructure network with no isolation or VLAN isolation, a load balancer backend and internet facing network with PVLAN, and tenant networks with isolation using network virtualization.
 - You can only use one type of isolation on a single logical network. If you do need this you'll need multiple logical networks.

- There's a practical limit of ~2000 tenants and ~4000 VM networks for a single VMM server.
- **Network sites:** If your organization has different locations and datacenters you can include information about those sites in your logical network settings. For example you could specify a New York site with and IP subnet and/or VLAN settings, and then a London site with different IP/VLAN settings. You can then assign IP address to VMs based on network, location, and VLAN settings. Note that:
 - You need to assign an IP subnet to a site if VMM is going to distribute static IP addresses to VMs in the site. If you're using DHCP you don't need a subnet.
 - You need to configure a VLAN if one's used in your physical site. If you're not using VLANs and you're using DHCP you don't need to define network sites in your logical network.

Create logical networks automatically

If you want VMM to automatically create logical networks (and VM networks) you can specify how VMM determines the logical network name.

1. Click **Settings > General**. Double-click **Network Settings**.
2. Configure the **Logical network matching** setting. Note that:
 - For Hyper-V hosts you can use the entire DNS suffix label, or the first one. For example if the DNS suffix is corp.contoso.com the logical network will be corp-contoso.com or just corp. This isn't supported for VMware hosts.
 - For Hyper-V and VMware hosts you can select the network connection name or the virtual network switch name (the name of the virtual network switch to which the physical adapter of the host is bound).
 - By default VMware hosts use the virtual network switch option.
 - You can also specify a fallback option if the first logical matching fails.

If you don't want VMM to create logical and VM networks automatically you can disable the global setting.

1. Click **Settings > General** and double-click **Network Settings**.
2. Clear **Create logical networks automatically**.

Create logical networks manually

1. **Fabric > Home > Show > Fabric Resources**. In **Fabric** expand **Networking > Logical Networks > Home > Create > Create Logical Network**.
2. In **Create Logical Network Wizard > Name** specify a name and description.
3. Specify how you want to isolate VM networks associated with this logical network:
 - If you want to have multiple VM networks associated with the logical network and isolate them using network virtualization click **One connected network > Allow new VM networks created on this logical network to use network virtualization**.
 - If you want to have multiple VM networks associated with the logical network and isolate them using a VLAN/PVLAN select **VLAN-based independent networks** or **Private VLAN (PVLAN) networks**.
 - If you don't want to isolate networks in the logical network, click **One connected network > Create a VM network with the same name to allow virtual machines to access this logical network directly**. With this setting you'll have a single VM network associated with your logical network.
 - If you've deployed a Microsoft network Controller in the VMM fabric you can select to have the logical network managed by the network controller.
4. In **Network Site** add network sites to the logical network. If you don't need to create network sites click

Next.

- **DHCP no VLAN:** If you're using DHCP to allocate IP addresses and you don't have VLANs you don't need a network site. Note that VMM automatically suggests a site name. Any network name shouldn't be longer than 64 characters.
 - **Static IP:** If you're using static IP addresses create at least one network site and associate at least one IP subnet with it.
 - **VLAN:** If you're using VLANs with static IP addressing create corresponding network site for the VLAN and subnet pairs. If you're using DHCP create corresponding network sites for VLAN information only.
 - **Network virtualization:** If you're using network virtualization create at least one network site with an associated IP subnet so that the logical network has an IP address pool.
 - **Load balancer:** If the logical network will contain a load balancer create at least one network site with an associated IP subnet.
5. If you're using an external network managed by a vendor network management console or virtual switch extension manager outside VMM you can configure settings in the vendor console and import them into VMM.
 6. In **Host groups that can use this network site** select each host group to which you want to make the logical network available.
 7. In **Associated VLANs and IP subnets** click **Insert Row** to specify the settings that you want to assign to the network site. If you selecting PVLAN you'll need to add a **SecondaryVLAN** for each VLAN. Ensure that the VLANs and subnets are available in your physical network. If you leave the VLAN field empty VMM assigns a value of 0 to indicate that VLANs aren't used. In trunk mode 0 indicates native VLAN.
 8. In **Summary** review the settings and click **Finish**. When the job shows as **Completed** verify the logical network in **Logical Networks and IP Pools**.

Next steps

If you created network sites and associated one or more IP subnets with them (you're not using DHCP) you can create static IP address pools from those subnets. Then VMM can automatically allocate IP addresses to VMs in the network site. [Set up IP address pools](#).

Set up VM networks in the VMM fabric

12/7/2018 • 4 minutes to read

This article describes how to create VM networks based on System Center - Virtual Machine Manager (VMM) logical networks.

In a virtualized network environment, we want to abstract virtual machines from the underlying logical network. VM networks help you to do this. VM networks are abstract objects that act as an interface to logical networks.

- A logical network can have one or more associated VM networks.
- If a logical network isn't isolated, then a single VM network will be associated with it.
- If a logical network is isolated, then multiple VM networks can be associated with it. These multiple VM networks allow us to use networks for different purposes. For example, as a provider you might want to host workload for multiple tenants on a single logical network, using a separate VM network for each tenant. The type of VM network you set up depends on the isolation settings for the logical network:
 - **Network virtualization:** If the logical network is isolated using network virtualization you can create multiple VM networks for a logical network. Within a VM network tenants can use any IP addresses they want for their VMs regardless of the IP addresses used on other VM networks. Tenants can also configure some network settings.
 - **VLAN:** If the logical network is isolated using VLAN or PVLAN you'll create one VM network for each network site and VLAN in the logical network.
 - **No isolation:** If the logical network is configured without isolation you'll create a single VM network linked to a logical network.

Before you start

- In some circumstances you'll need to create a static IP address pool on the VM network after you've created it. [Learn more.](#)
- By default machines within a specific VM network can connect to each other. If your VM network will connect to other networks you can configure it with a gateway (network service). If you want to add a gateway to the VM network you'll need to create it. [Learn more.](#)

Create a VM network (network virtualization)

1. Click **VMs and Services > VM Networks > Home > Create > Create VM Network**.
2. In **Create VM Network Wizard > Name**, type in a name and description and select a logical network on which to base the VM network.
3. In **Isolation**, select **Isolate using Hyper-V network virtualization**, and verify the IP address protocols.
4. In **VM Subnets** click **Add**, and specify subnets for the VM network using CIDR notation. You can add multiple subnets.
5. In **Connectivity**, if you see the message **No network service**, it specifies a gateway has been added to VMM and you can click **Next**. If you don't see the message, specify the gateway (network service) options:
 - **No connectivity:** Leave all check boxes cleared if the virtual machines on this VM network will communicate only with other virtual machines on this VM network. You can also leave clear if you plan to configure the gateway later.

- **Connect to another network through a VPN tunnel:** Select this option if the virtual machines on this VM network will communicate with other networks over VPN. If the device will use the Border Gateway Protocol, enable this protocol. Select the VPN gateway device that you want to use. Confirm the settings. If the **VPN Connections** or Border Gateway Protocol pages appear, complete the settings based on information from the VPN admin. page appears. If you selected the check box for Border Gateway Protocol, the Border Gateway Protocol page also appears.
- **Connect directly to an additional logical network:** Select this option if the virtual machines on this VM network will communicate with other networks in this data center. Select either direct routing or NAT. Select the gateway device you want to use and confirm the settings.

6. In **Summary** verify settings and click **Finish**. Verify the network in **VM Networks and IP Pools**.

Create a VM network (VLAN/PVLAN)

1. Click **VMs and Services > VM Networks > Home > Create > Create VM Network**.
2. In **Create VM Network Wizard > Name** type in a name and description and select a logical network on which to base the VM network.
3. In **Isolation Options**:
 - Select **Automatic** if you want VMM to automatically configure the isolation settings for the VM network. VMM will select a network site and subnet VLAN based on those available in the logical network.
 - Select **Specify a VLAN** to configure isolation manually. Note that tenant administrators can only select the **Automatic** option.
4. In **Summary** verify settings and click **Finish**. Verify the network in **VM Networks and IP Pools**.

Create a VM network (no isolation)

1. Click **VMs and Services > VM Networks > Home > Create > Create VM Network**.
2. In **Create VM Network Wizard > Name** type in a name and description and select a logical network on which to base the VM network.
3. In **VM Networks and IP Pools** check for a VM network with the same name as the logical network you want to give direct access to. If one exists it probably indicates that the VM network was created automatically when you created the logical network. You can check whether the VM network provides direct access by clicking its properties. If **Name** and **Access** are the only tabs it provides direct access.
4. If there's no existing VM network click **Home > Create > Create VM Network**.
5. In **Summary** verify settings and click **Finish**. Verify the network in **VM Networks and IP Pools**.

Set up static IP address pools in the VMM fabric

6/14/2019 • 5 minutes to read

This article describes how to set up static IP address pools for logical and VM networks in the System Center - Virtual Machine Manager (VMM) networking fabric.

When you set up the logical network you'll need to configure a static IP address pool if you're not using DHCP. In some circumstances you'll need to create IP address pools on the logical network only, and in others you'll need to create the pool on both the logical and VM networks:

- **Pool on logical and VM network:** If you configure your logical network for network virtualization you'll need to create IP address pools on the logical network and the VM network.
- **Pool on logical network only:** If you're using VLAN or no isolation you can use DHCP or create IP address pools on the logical network only. They'll automatically become available on the VM network.
- **Imported address pools:** If you are using external networks through a vendor console your IP address pools will be imported from the vendor and you don't need to create them in VMM.

Create a static address pool for a logical network

1. In **Logical Networks and IP Pools** click the logical network > **Home** > **Create** > **Create IP Pool**.
2. In **Create Static IP Address Pool Wizard** > **Name** specify a name and description. Make sure the correct logical network is indicated.
3. In **Network Site** select to use an existing site and select the IP subnet, or create a new site.
 - For an existing site select the site and IP subnet from which to create the pool.
 - For a new site specify the site name, IP subnet to assign to the site, and VLAN information if relevant. Select the host groups that can access this site and the logical network.
4. If you're using network virtualization you can use the pool to support multicasting or broadcasting. To do this click **Create a multicast IP address pool** and select the IP subnet you want to use. To use multicasting or broadcasting note that:
 - The logical network must have network virtualization enabled.
 - The IP protocol setting for the VM network must match the IP protocol settings for the underlying logical network. You can't view the protocol setting in the VMM console after you've created it. You'll need to run **Get-SCVMMNetwork -Name <VM network name> | Format -List Name, Isolation Type, PoolType** to see it.
 - After you've configured this feature multicast and broadcast packets on the VM network will use the IP addresses from the multicast IP address pool. Each subnet in the VM network will consume one IP address from the multicast pool.
5. In **IP address range** enter the start and end address for the pool. They must be contained within the subnet. In **VIPs and reserved IP addresses** specify IP address range you want to reserve for VIPs. VIPs are used during deployment of a service in a load-balanced service tier. VMM automatically assigns a VIP to the load balancer from the reserved VIP address range.
6. In **Gateway** click **Insert** if you want to specify one or more default gateways and the metric. The default gateway address must be in the same subnet range as the IP address pool but doesn't need to be part of the pool.
7. In **DNS** specify DNS information, including DNS servers, the default DNS suffix for the connection, and

the list of DNS search suffixes.

8. In **WINS** click **Insert** if you want to enter the IP address of a WINS server. You can also select whether to enable NetBIOS over TCP/IP. This isn't recommended if the address range is made up of public addresses.
9. In **Summary** verify the settings and click **Finish**. When the job shows as **Completed** verify the pool in **Logical Networks and IP Pools**.

Set up an IP address pool on a VM network

1. Click **VMs and Services > VM Networks > Home > Show > VM Networks > VM Network** tab.
2. In **VM Networks and IP Pools** click the VM network **Create > Create IP Pool**.
3. In **Create Static IP Address Pool Wizard > Name** specify a name and description. Make sure the correct logical network is indicated. Check that the correct VM network and subnet is selected.
4. In **IP address range** enter the start and end address for the pool. You can create multiple IP address pools in a subnet but the ranges mustn't overlap. In **Reserved IP addresses** specify any ranges you want to reserve for other purposes.
5. In **Gateway** click **Insert** if you want to specify one or more default gateways and the metric. The default gateway address must be in the same subnet range as the IP address pool but doesn't need to be part of the pool.
6. In **DNS** specify DNS information, including DNS servers, the default DNS suffix for the connection, and the list of DNS search suffixes. For virtual machines that will join an Active Directory domain, we recommend that you use Group Policy to set the primary DNS suffix. This will ensure that when a Windows-based virtual machine is set to register its IP addresses with the primary DNS suffix, a Windows-based DNS server will register the IP address dynamically. Additionally, the use of Group Policy enables you to have an IP address pool that spans multiple domains. In this case, you would not want to specify a single primary DNS suffix.
7. In **WINS** click **Insert** if you want to enter the IP address of a WINS server. You can also select whether to enable NetBIOS over TCP/IP. This isn't recommended if the address range is made up of public addresses.
8. In **Summary** verify the settings and click **Finish**. When the job shows as **Completed** verify the pool in **Logical Networks and IP Pools**.

Release inactive addresses from the static address pool

You can release inactive addresses. When you do VMM returns the address to the static IP or MAC address pool and considers it available for reassignment. An address is considered inactive if:

- A host that was assigned a static IP address through the bare-metal deployment process is removed from VMM management. When you remove the host, any IP and MAC addresses that were statically assigned to virtual machines on the host are also marked as inactive.
- A virtual machine goes into a missing state because it was removed outside VMM.

1. Release the IP addresses:
 - To release addresses in a pool in a logical network click **Logical Networks and IP Pools** expand the logical network and click the IP address pool.
 - To release addresses in a pool in a VM network click **Logical Networks and IP Pools** expand the VM network and click the IP address pool.
2. Click **Home > Properties > Inactive addresses** and select the inactive IP addresses that you want to release.

Next steps

[Create a VM network](#)

Add a network virtualization gateway to the VMM fabric

3/14/2019 • 4 minutes to read

Read this article to learn about setting up network virtualization gateways in the System Center - Virtual Machine Manager (VMM) networking fabric.

By default, if you're using isolated VM networks in your VMM fabric, VMs associated with a network can only connect to machines in the same subnet. If you want to connect VMs further than the subnet, you'll need a gateway.

Network virtualization

You set up network virtualization so that multiple VM networks are overlaid on the VMM logical networks that model your physical network topology and thus decouple the VM networks from the physical network infrastructure. Network virtualization uses NVGRE (Network Virtualization using Generic Routing Encapsulation) to virtualize IP addresses. Review the following to learn more about [NVGRE](#).

To figure out whether you need a network virtualization gateway in your network consider:

- Do you need to connect from VMs in isolated VM networks to other on-premises apps?
- Do you need to connect from isolated VMs to the internet?
- Do you need to connect from isolated VM networks to shared services such as DNS?

You can set up your gateway in a number of ways depending on your requirements:

- Connectivity to a public network can be achieved through NAT.
- Connectivity to an on-premises network is over a VPN tunnel (with or without Border Gateway Protocol (BGP))
- Direct routing without NAT can be used for connectivity between different VM networks.

Prerequisites

- **Provider software:** If you want to use a non-Windows gateway device, you'll need the provider and an account with permissions to configure the gateway. You install the provider on the VMM server. If certificates are required (for example if the gateway is in an untrusted domain) you'll need to be able to view thumbprint information for those certificates.
- **Windows Server gateway:** If you want to configure a gateway running Windows Server, you can use a predefined template available from the Microsoft Download Center. The template supports System Center 2012 R2 or later versions.
- **Logical networks:** You need logical networks (you'll need more than one if you want the gateway to connect from VM networks in one logical network to VM networks in another).
- **Remote VPN settings:** If you want to connect the gateway to a remote VPN server you'll need:
 - The remote server IP address and information about on-premises subnets or the BGP address if relevant.
 - You'll need to identify how you'll authenticate with the remote VPN server. If it uses a preshared key you can authenticate with a Run As account and specify the shared key as the password. Or you can authenticate with a certificate. The certificate can be either a certificate that the remote VPN server selects automatically or a certificate that you have obtained and placed on your network.
 - Check whether you need specific VPN connection settings (encryption, integrity checks, cipher

transforms, authentication transforms, Perfect Forward Secrecy (PFS) group, Diffie-Hellman group, and VPN protocol) or you can use the default settings.

Add a Windows Server Gateway

The service template provides a highly available Windows Server Gateway deployment in active-standby mode.

1. You'll need to download the template from the [Download center](#).

NOTE

The templates downloaded are applicable for VMM 2012 R2, 2016 and 1801.

2. The download is a compressed zip file. You'll need to extract the file. Files include a user guide, two service templates, and a custom resource folder (a folder with a .cr extension) that contains files required for the service templates.
3. You'll need to decide which template to use, and then follow the instructions in the Quick Start Guide. The guide includes prerequisites for the template deployment, and instructions for setting up logical networks, creating a scale-out file server, preparing virtual hard disks for the gateway VM, and copying the custom resource file to the library. After you've set up the infrastructure it describes how to import and customize the template, and how to deploy it. There's also troubleshooting information if issues arise.

Add a non-Windows gateway

NOTE

You'll need to install the provider software on the VMM management server and add the gateway to the fabric. Obtain the provider software. You can review a list of supported providers in **Settings > Configuration Providers**

Use the following procedure to add the non-windows gateway:

1. Click **Fabric > Network Service**. Right-click and select **Add Network Service** to open the Network Service wizard. Network services include gateways, virtual switch extensions, network managers, and top-of-rack (TOR) switches. or On Home, click **Add Resources > Network Service**.
2. In **Add Network Service Wizard > Name** specify a name and description for the gateway.
3. In **Manufacturer and Model** click the required settings.
4. In **Credentials**, specify a Run As account with permissions in the domain to which the gateway is connected.
5. In **Connection String** type the string that the gateway should use. The string syntax is defined by the gateway vendor.
6. In **Certificates** if listed, verify the thumbprints of the certificates match those installed on the gateway. Select to confirm that the certificates can be imported. If none is listed the gateway probably doesn't need certificate authentication. If they're needed make sure they're installed correctly on the gateway.
7. In **Gather Information** click **Scan Provider** to run the basic validation test against the gateway.
8. In **Host Group** select one or more host groups to which the gateway will be available.
9. In **Summary**, review the settings and click **Finish**.
10. After the gateway is added find its listing in **Network Services** and right-click > **Properties > Connectivity**.
11. Select **Enable front end connection** and select the gateway network adapter and network site that provides connectivity outside the enterprise datacenter or hosting provider. Select **Enable back end connection** and select a gateway network adapter and network site in a logical network within the enterprise. The network must have network virtualization enabled and the network site must have a static IP address.
12. When you create a VM network you can assign the gateway to it, and select the required connectivity options.

Set up port profiles in the VMM fabric

12/7/2018 • 9 minutes to read

Use this article to learn about and set up uplink port profiles and virtual network adapter port profiles in the System Center Virtual Machine Manager (VMM) networking fabric.

- **Uplink port profiles:** You create uplink port profiles and then apply them to physical network adapters when you deploy switches. Uplink port profiles define the load balancing algorithm for an adapter, and specify how to team multiple network adapters on a host that use the same uplink port profile. This profile is used in conjunction with the logical network that you associated with the adapter.
- **Virtual network adapter port profiles.** You apply virtual network adapter port profiles to virtual network adapters. These profiles define specific capabilities, such as bandwidth limitations, and priority. VMM includes a number of built-in profiles.
- **Port classifications:** After creating a virtual network adapter port profiles you can create port classifications. Port classifications are abstractions that identify different types of virtual network adapter port profiles. For example you could create a classification called FAST to identify ports that are configured to have more bandwidth and another one called SLOW with less bandwidth. Classifications are included in logical switches. Administrators and tenants can choose a classification for their VM virtual machine adapters. By default, VMM includes built-in classifications that map to the built-in virtual network adapter port profiles

Define uplink port profiles

Some guidelines for creating uplink port profiles:

- You need at least one uplink port profile for each physical network in your environment. If you do have a simple environment with a single physical network and all hosts are configured the same way, with the same protocols for network adapter teaming, then you might only need a single uplink port profile. This is rare though. You'll probably need to scope or restrict certain logical networks to a specific group of hosts computer, and this need makes it useful to create multiple uplink port profiles.
- You need to define uplinks for each physical location that has its own VLAN and IP subnets.
- If you plan to restrict or otherwise scope logical networks to a specific set of host computers, you will need to create uplinks for each group of computers.
- You need separate uplink port profiles for groups of computers (in each physical location) that have different connectivity requirements or use different teaming protocols.
- You might consider creating separate uplinks for networks that do not or will not support network virtualization.
- Network sites that will be included in a profile should be scoped to the same group of host computers. If they aren't you'll receive an out-of-scope error when you try to apply it to a computer that isn't a member of the host groups defines in every one of the network sites included in the uplink profile.
- You should try and ensure that each of the network sites that you add to an uplink port profile refers to a different logical network. If you do otherwise all of the VLANs and IP subnets defined in those network sites will be associated with the logical network on any host computer on which the uplink port profile is applied. If you're not using VLAN isolation the host computer has no way to establish which of the range of possible VLANs and IP subnets will be needed to allow VMs connected to the logical network
- You can create an uplink port profile that contains references to multiple network sites (and hence logical networks). You should ensure that the VLANs and IP addresses in each of the selected sites should be valid (routable) from the physical network adapter to which the port profile has been applied.
- When you apply the profile to a physical network adapter these network sites determine the set of logical

networks that should be associated with the physical adapter and the VLANs and IP subnets that will be allocated to VMs and services that connect to those logical networks.

Create an uplink port profile

1. Click **Fabric > Home > Show > Fabric Resources**. Click **Fabric** tab > **Networking > Port Profiles > Hyper-V Port Profile**.
2. In **Create Hyper-V Port Profile Wizard > General** type in a name, description and select **Uplink Port Profile**. Select the load balancing and teaming settings:
 - **Load balancing: Host Default** is the default setting and this will either distribute network traffic based on the Hyper-V switch port identifier of the source VM or use a **Dynamic** loading balancing algorithm, depending what the Hyper-V host supports. Dynamic is only available from Windows Server 2012 R2 onwards. You can also select:
 - Hyper-V port: Distributes network traffic according to the Hyper-V switch port identifier of the source VM.
 - Transport ports: Uses the source and destination TCP ports and the IP addresses to create a hash and then assigns the packets that have that hash value to one of the available network adapters.
 - IP addresses: Uses the source and destination IP addresses to create a hash and then assigns the packets that have that hash value to one of the available network adapters.
 - MAC addresses: Uses the source and destination MAC addresses to create a hash and then assigns the packets that have that hash value to one of the available network adapters.
 - **Teaming: Switch Independent** is the default setting and this specifies that physical network switch configuration isn't needed for the NIC team. The network switch is not configured and so allow network adapters within the team to be connected to multiple (non-trunked) physical switches. You can also select:
 - **LACP**: Use the LACP protocol to dynamically identify links that are connected between the host and a given switch.
 - **Static teaming**: You configure both the switch and host to identify which links form the team.
3. In **Network Configuration** select one or more network sites for this uplink port profile to support. Uplink port profiles contains a list of network sites with each network site representing a link to a different logical network. Select **Enable Hyper-V Network Virtualization** if you want to enable network virtualization to deploy multiple VM networks on a single physical network. You should only do this if the logical network is configured for network virtualization with **Allow new VM networks created on this logical network to use network virtualization** enabled.
4. In **Summary** review the settings and click **Finish**.

After you create an uplink port profile, the next step is to add it to a logical switch, which places it in a list of profiles that are available through that logical switch. When you apply the logical switch to a network adapter in a host, the uplink port profile is available in the list of profiles, but it is not applied to that network adapter until you select it from the list. This helps you to create consistency in the configurations of network adapters across multiple hosts, but also enables you to configure each network adapter according to your specific requirements.

Create a virtual network adapter port profile

1. Click **Fabric > Home > Show > Fabric Resources**. Click **Fabric** tab > **Networking > Port Profiles > Home > Create > Hyper-V Port Profile**.
2. In **Create Hyper-V Port Profile Wizard > General** type in a name, description and select **Uplink Port Profile**.
3. In **Offload Setting**, specify a setting for offloading traffic:

- **Enable virtual machine queue (VMQ):** Packets destined for a virtual network adapter are delivered directly to a queue for that adapter, and they do not have to be copied from the management operating system to the virtual machine. The physical network adapter must support VMQ.
- **Enable IPsec task offloading:** Some or all of the IPsec computational work is shifted from the computer's CPU to a dedicated processor on the network adapter. The physical network adapter and the guest operating system must support it.
- **Enable single-root I/O virtualization:** A network adapter can be assigned directly to a virtual machine. This maximizes network throughput while minimizing network latency and minimizing the CPU overhead that is required to process network traffic. The physical network adapter and drivers in the management operating system and guest operating system must support it. If you want to use SR-IOV you'll need to enable it in the port profile (in **Offload** settings) and in the logical switch (**General** settings) that includes the port profile. It must be configured correctly on the host when you create the virtual switch that brings port settings and the logical switch you want to use on the host together. In the virtual switch you attach the port profile to the virtual switch using a port classification (either the default SR-IOV classification provided by VMM, or a custom one)

4. In **Security Settings** specify:

- **Allow MAC spoofing:** Allows a virtual machine to change the source MAC address in outgoing packets to an address that is not assigned to that virtual machine. For example, a load-balancer virtual appliance might require this setting to be enabled.
- **Enable DHCP guard:** Helps protect against a malicious virtual machine that represents itself as a DHCP server for man-in-the-middle attacks.
- **Allow router guard:** Helps protect against advertisement and redirection messages that are sent by an unauthorized virtual machine that represents itself as a router.
- **Allow guest teaming:** Allows you to team the virtual network adapter with other network adapters that are connected to the same switch.
- **Allow IEEE priority tagging:** Allows you to tag outgoing packets from the virtual network adapter with IEEE 802.1p priority. These priority tags can be used by Quality of Service (QoS) to prioritize traffic. If IEEE priority tagging is not allowed, the priority value in the packet is reset to 0.
- **Allow guest specified IP addresses:** Affects VM networks using network virtualization. The VM (guest) can add and remove IP addresses on this virtual network adapter. This can simplify the process of managing virtual machine settings. Guest-specified IP addresses are required for virtual machines that use guest clustering with network virtualization. The IP address that a guest adds must be within an existing IP subnet in the VM network.

5. In **Bandwidth Settings** specify the minimum and maximum bandwidth that are available to the adapter. The minimum bandwidth can be expressed as megabits per second (Mbps) or as a weighted value (from 0 to 100) that controls how much bandwidth the virtual network adapter can use in relation to other virtual network adapters. Note that bandwidth settings aren't used if SR-IOV is enabled on the port profile and logical switch that contains the port profile.

6. In **Summary** review the settings and click **Finish**.

After creating a port profile you can create a port classification.

Create port classifications for virtual network adapter port profiles

1. Click **Fabric > Home > Show > Fabric Resources**. Click the **Fabric** tab > **Networking > Port Classifications > Home > Create > Port Classification**.
2. In **Create Port Classification Wizard > Name** specify a classification name.

Next steps

Create logical switches

5/16/2019 • 22 minutes to read

This article describes how to create logical switches in the System Center - Virtual Machine Manager (VMM) fabric.

How to create logical switches

This article describes how to create logical switches in the System Center - Virtual Machine Manager (VMM) fabric, convert a host virtual switch to a logical switch, and set up virtual switch extensions if you need them.

A logical switch brings virtual switch extensions, port profiles, and port classifications together so that you can configure each network adapter with the settings you need, and have consistent settings on network adapters across multiple hosts. You can team multiple network adapters by applying the same logical switch and uplink port profile to them.

Set up virtual switch extensions

You install switch extensions on the VMM server and then include them in a logical switch. There are a few types of switch extensions:

- **Monitoring extensions** can be used to monitor and report on network traffic, but they cannot modify packets.
- **Capturing extensions** can be used to inspect and sample traffic, but they cannot modify packets.
- **Filtering extensions** can be used to block, modify, or defragment packets. They can also block ports.
- **Forwarding extensions** can be used to direct traffic by defining destinations, and they can capture and filter traffic. To avoid conflicts, only one forwarding extension can be active on a logical switch.

You can set up a virtual switch extension manager (network manager) if you want to managed extensions using a vendor management console and the VMM console together.

Set up a virtual switch extension manager

1. Obtain the provider software from your vendor and install the provider on the VMM management server. If you have a cluster install it on all nodes.
2. Click **Fabric > Home > Show > Fabric Resources > Networking > Switch Extension Managers**.
3. In **Add Virtual Switch Extension Manager Wizard > General** specify the manufacturer and type the connection string for example myextmanager1.contoso.com:443. The exact syntax is defined by the vendor. Specify the account you want to use to connect to the resource.
4. In **Host Groups** specify the host groups for which you want to use the extension manager.
5. In **Summary** review settings and click **Finish**. Check that the extension appears in the **Virtual Switch Extension Managers** pane.

Set up a logical switch

1. Make sure you have at least one uplink port profile before you begin.
2. Click **Fabric** tab > **Networking > Logical Switches > Create Logical Switch**.
3. In **Create Logical Switch Wizard > Getting Started**, review the information.
4. In **General**, specify a name and optional description.

5. In **Uplink Mode**, select:
 - **No Uplink Team** if you're not using teaming.
 - **Embedded Team** if you want to deploy the switch with SET-based teaming
 - **Team** if you want to use NIC teaming
6. In **Settings**, select the minimum bandwidth mode. If you've deployed Microsoft network controller, you can specify that it should manage the switch. If you enable this setting you won't be able to add extensions to the switch.
7. Enable SR-IOV if you need to. SR-IOV enables virtual machines to bypass the switch and directly address the physical network adapter. If you want to enable:
 - Make sure that you have SR-IOV support in the host hardware and firmware, the physical network adapter, and drivers in the management operating system and in the guest operating system.
 - Create a native port profile for virtual network adapters that is also SR-IOV enabled.
 - When you configure networking settings on the host (in the host property called Virtual switches), attach the native port profile for virtual network adapters to the virtual switch by using a port classification. You can use the SR-IOV port classification that is provided in VMM, or create your own port classification.
8. In **Extensions**, if you're using virtual switch extensions select them and arrange the order. extensions process network traffic through the switch in the order you specify. Note that only one forwarding extension can be enabled.
9. In **Virtual Port** add one or more port classifications and virtual network adapter port profiles. You can also create a port classification and set a default classification.
10. In **Uplink** add an uplink port profile, or [create a new one](#). When you add an uplink port profile, it is placed in a list of profiles that are available through that logical switch. However, when you apply the logical switch to a network adapter in a host, the uplink port profile is applied to that network adapter only if you select it from the list of available profiles.
11. In **Summary** review the settings and click **Finish**. Verify the switch appears in **Logical Switches**.

Convert virtual switch to logical switch

If a host in the VMM fabric has a standard virtual switch, you can convert it to use as a logical switch.

NOTE

- The following procedure is not applicable for SET, use the [script](#) instead.
- Before you can convert, you need a logical switch in place, with specific settings.
- You must be a member of the Administrator user role, or a member of the Delegated Administrator user role, where the management scope includes the host group in which the Hyper-V host is located.

Compare switch settings

1. In **Server Manager** on the host, click **Hyper-V**. Close Server Manager.
2. Right-click the host > **Configure NIC Teaming**, and record any teaming and load balancing settings.
3. In **Hyper-V Manager**, right-click the host > **Virtual Switch Manager**. Select the virtual switch and verify whether **Enable single-root I/O virtualization (SR-IOV)** is selected. Close Hyper-V Manager.
4. In the VMM console > **Fabric** > **Servers** > **All Hosts**, right-click the host > **Properties**.
5. In **Virtual Switches**, note the properties, including logical network, and minimum bandwidth mode.

6. In **Fabric > Networking > Logical Switches**, right-click the logical switch that you want to convert the host configuration to, and click **Properties**.
7. In **Logical Switches**, record the information:
 - In **General**, record the uplink mode, whether SR-IOV is enable, and minimum bandwidth mode.
 - In **Extensions**, note whether any forwarding extensions have been added to the logical switch.
 - In **Virtual port**, record the names of the port profiles that are listed. Be sure to note if one of them has SR-IOV in the name.
 - In **Uplinks**, record the network sites, whether uplink mode is teamed, the load balancing algorithm, and teaming mode.
8. In **Fabric > Networking**, click **Port Profiles**. For any relevant port profiles, click **Properties**. In **Offload Settings**, see if **Enable Single-root I/O virtualization** is checked.
9. Now compare the recorded information that you recorded for the logical switch and port profiles, with the virtual switch information.
10. Review the following table to see whether you can convert the host to use the logical switch.

ITEM	CONVERSION
SR-IOV	<p>The SR-IOV setting (enabled or disabled) must be the same in the logical switch as it is in the virtual switch.</p> <p>If SR-IOV is enabled, it must be enabled in the logical switch itself, and in at least one virtual network adapter port profile within the logical switch.</p>
Uplink mode Load balancing algorithm Teaming mode	<p>The Uplink mode setting must match.</p> <p>If the uplink mode is Team, then the Load balancing algorithm and Teaming mode must also match.</p>
Minimum bandwidth mode	Must match.
Network sites	The logical switch must be configured for the correct network sites (in the correct logical network) for this host.

11. If the settings in the logical switch don't match as described in the table, you need to find or create a logical switch that does match.

Convert a host to use a logical switch

NOTE

- The conversion will not interrupt network traffic.
- If any operation in the conversion fails, no settings will be changed, and the switch will not be converted.

1. In VMM, click **Fabric > Servers > All Hosts**. Right-click the host > **Properties**.
2. On the **Virtual Switches** tab, click **Convert to Logical Switch**.
3. Select the logical switch that you want to convert the host to. Then select the uplink port profile to use, and click **Convert**.
4. The **Jobs** dialog box might appear, depending on your settings. Make sure that the job has a status of **Completed**, and then close the dialog box.
5. To verify that the switch was converted, right-click the host, click **Properties**, and then click the **Virtual**

Switches tab.

Script for SET switch conversion

NOTE

Create a logical switch in VMM with the same name as the SET switch that is deployed on the host. Standard switch will be converted to this logical switch after you run the following script on the host.

```
#Replace Virtual Switch name with already deployed switch name on host
$VirtualSwitchName="SETswitch"

#Replace logical switch ID below with the one got from Get-SCLogicalSwitch cmdlet for the switch created in VMM
$LogicalSwitchId="45b98a8d-1887-4431-9f20-8b9beed853ce"

#Replace the port profile set name with the one created and associated with the above logical switch in VMM
$PortProfileSetName="Mgmt_UPP"

#Replace uplink port profile set ID with the one got from Get-SCUplinkPortProfileSet for the port profile set created in VMM
$PortProfileSetId="fd9e4c9a-4ffa-4845-808d-930e6616b62f"

$vswitch=Get-VMSwitch -Name $VirtualSwitchName
$VMMPortFeatureId="1f59a509-a6ba-4aba-8504-b29d542d44bb"
$defaultPortFeature = Get-VMSwitchExtensionPortFeature -FeatureId $VMMPortFeatureId
$VMMFeatureId="8b54c928-eb03-4aff-8039-99171dd900ff"
$currentFeature = Get-VMSwitchExtensionSwitchFeature -SwitchName $VirtualSwitchName -FeatureId $VMMFeatureId
$defaultFeature = Get-VMSwitchExtensionSwitchFeature -FeatureId $VMMFeatureId
$defaultFeature.SettingData.LogicalSwitchId=$LogicalSwitchId
$defaultFeature.SettingData.LogicalSwitchName=$VirtualSwitchName
Add-VMSwitchExtensionSwitchFeature -SwitchName $VirtualSwitchName -VMSwitchExtensionFeature $defaultFeature

$defaultPortFeature = Get-VMSwitchExtensionPortFeature -FeatureId $VMMPortFeatureId
$defaultPortFeature.SettingData.PortProfileSetId=$PortProfileSetId
$defaultPortFeature.SettingData.PortProfileSetName=$PortProfileName
$defaultPortFeature.SettingData.NetCfgInstanceId="{ " + $vswitch.Id + "}"
Add-VMSwitchExtensionPortFeature -SwitchName $VirtualSwitchName -VMSwitchExtensionFeature $defaultPortFeature
-ExternalPort
```

After you run the script, refresh the host in VMM and verify if VMM recognizes the switch as logical switch.

How to create logical switches

This article describes how to create logical switches in the System Center - Virtual Machine Manager (VMM) fabric, convert a host virtual switch to a logical switch, and set up virtual switch extensions if you need them.

A logical switch brings virtual switch extensions, port profiles, and port classifications together so that you can configure each network adapter with the settings you need, and have consistent settings on network adapters across multiple hosts. You can team multiple network adapters by applying the same logical switch and uplink port profile to them.

Set up virtual switch extensions

You install switch extensions on the VMM server and then include them in a logical switch. There are a few types of switch extensions:

- **Monitoring extensions** can be used to monitor and report on network traffic, but they cannot modify packets.
- **Capturing extensions** can be used to inspect and sample traffic, but they cannot modify packets.

- **Filtering extensions** can be used to block, modify, or defragment packets. They can also block ports.
- **Forwarding extensions** can be used to direct traffic by defining destinations, and they can capture and filter traffic. To avoid conflicts, only one forwarding extension can be active on a logical switch.

You can set up a virtual switch extension manager (network manager) if you want to managed extensions using a vendor management console and the VMM console together.

Set up a virtual switch extension manager

1. Obtain the provider software from your vendor and install the provider on the VMM management server. If you have a cluster install it on all nodes.
2. Click **Fabric > Home > Show > Fabric Resources > Networking > Switch Extension Managers**.
3. In **Add Virtual Switch Extension Manager Wizard > General** specify the manufacturer and type the connection string for example myextmanager1.contoso.com:443. The exact syntax is defined by the vendor. Specify the account you want to use to connect to the resource.
4. In **Host Groups** specify the host groups for which you want to use the extension manager.
5. In **Summary** review settings and click **Finish**. Check that the extension appears in the **Virtual Switch Extension Managers** pane.

Set up a logical switch

1. Make sure you have at least one uplink port profile before you begin.
2. Click **Fabric** tab > **Networking > Logical Switches > Create Logical Switch**.
3. In **Create Logical Switch Wizard > Getting Started**, review the information.
4. In **General**, specify a name and optional description.
5. In **Uplink Mode**, select:
 - **No Uplink Team** if you're not using teaming.
 - **Embedded Team** if you want to deploy the switch with SET-based teaming
 - **Team** if you want to use NIC teaming
6. In **Settings**, select the minimum bandwidth mode. If you've deployed Microsoft network controller, you can specify that it should manage the switch. If you enable this setting you won't be able to add extensions to the switch.
7. Enable SR-IOV if you need to. SR-IOV enables virtual machines to bypass the switch and directly address the physical network adapter. If you want to enable:
 - Make sure that you have SR-IOV support in the host hardware and firmware, the physical network adapter, and drivers in the management operating system and in the guest operating system.
 - Create a native port profile for virtual network adapters that is also SR-IOV enabled.
 - When you configure networking settings on the host (in the host property called Virtual switches), attach the native port profile for virtual network adapters to the virtual switch by using a port classification. You can use the SR-IOV port classification that is provided in VMM, or create your own port classification.
8. In **Extensions**, if you're using virtual switch extensions select them and arrange the order. extensions process network traffic through the switch in the order you specify. Note that only one forwarding extension can be enabled.
9. In **Virtual Port** add one or more port classifications and virtual network adapter port profiles. You can also create a port classification and set a default classification.
10. In **Uplink** add an uplink port profile, or [create a new one](#). When you add an uplink port profile, it is placed

in a list of profiles that are available through that logical switch. However, when you apply the logical switch to a network adapter in a host, the uplink port profile is applied to that network adapter only if you select it from the list of available profiles.

11. In **Summary** review the settings and click **Finish**. Verify the switch appears in **Logical Switches**.

Convert virtual switch to logical switch

If a host in the VMM fabric has a standard virtual switch, you can convert it to use as a logical switch.

NOTE

- The following procedure is not applicable for SET, use the [script](#) instead.
- Before you can convert, you need a logical switch in place, with specific settings.
- You must be a member of the Administrator user role, or a member of the Delegated Administrator user role, where the management scope includes the host group in which the Hyper-V host is located.

Compare switch settings

1. In **Server Manager** on the host, click **Hyper-V**. Close Server Manager.
2. Right-click the host > **Configure NIC Teaming**, and record any teaming and load balancing settings.
3. In **Hyper-V Manager**, right-click the host > **Virtual Switch Manager**. Select the virtual switch and verify whether **Enable single-root I/O virtualization (SR-IOV)** is selected. Close Hyper-V Manager.
4. In the VMM console > **Fabric** > **Servers** > **All Hosts**, right-click the host > **Properties**.
5. In **Virtual Switches**, note the properties, including logical network, and minimum bandwidth mode.
6. In **Fabric** > **Networking** > **Logical Switches**, right-click the logical switch that you want to convert the host configuration to, and click **Properties**.
7. In **Logical Switches**, record the information:
 - In **General**, record the uplink mode, whether SR-IOV is enable, and minimum bandwidth mode.
 - In **Extensions**, note whether any forwarding extensions have been added to the logical switch.
 - In **Virtual port**, record the names of the port profiles that are listed. Be sure to note if one of them has SR-IOV in the name.
 - In **Uplinks**, record the network sites, whether uplink mode is teamed, the load balancing algorithm, and teaming mode.
8. In **Fabric** > **Networking**, click **Port Profiles**. For any relevant port profiles, click **Properties**. In **Offload Settings**, see if **Enable Single-root I/O virtualization** is checked.
9. Now compare the recorded information that you recorded for the logical switch and port profiles, with the virtual switch information.
10. Review the following table to see whether you can convert the host to use the logical switch.

ITEM	CONVERSION
SR-IOV	<p>The SR-IOV setting (enabled or disabled) must be the same in the logical switch as it is in the virtual switch.</p> <p>If SR-IOV is enabled, it must be enabled in the logical switch itself, and in at least one virtual network adapter port profile within the logical switch.</p>

ITEM	CONVERSION
Uplink mode	The Uplink mode setting must match.
Load balancing algorithm Teaming mode	If the uplink mode is Team , then the Load balancing algorithm and Teaming mode must also match.
Minimum bandwidth mode	Must match.
Network sites	The logical switch must be configured for the correct network sites (in the correct logical network) for this host.

11. If the settings in the logical switch don't match as described in the table, you need to find or create a logical switch that does match.

Convert a host to use a logical switch

NOTE

- The conversion will not interrupt network traffic.
- If any operation in the conversion fails, no settings will be changed, and the switch will not be converted.

1. In VMM, click **Fabric > Servers > All Hosts**. Right-click the host > **Properties**.
2. On the **Virtual Switches** tab, click **Convert to Logical Switch**.
3. Select the logical switch that you want to convert the host to. Then select the uplink port profile to use, and click **Convert**.
4. The **Jobs** dialog box might appear, depending on your settings. Make sure that the job has a status of **Completed**, and then close the dialog box.
5. To verify that the switch was converted, right-click the host, click **Properties**, and then click the **Virtual Switches** tab.

Script for SET switch conversion

NOTE

Create a logical switch in VMM with the same name as the SET switch that is deployed on the host. Standard switch will be converted to this logical switch after you run the following script on the host.

```
#Replace Virtual Switch name with already deployed switch name on host
$VirtualSwitchName="SETswitch"

#Replace logical switch ID below with the one got from Get-SCLogicalSwitch cmdlet for the switch created in VMM
$LogicalSwitchId="45b98a8d-1887-4431-9f20-8b9beed853ce"

#Replace the port profile set name with the one created and associated with the above logical switch in VMM
$PortProfileSetName="Mgmt_UPP"

#Replace uplink port profile set ID with the one got from Get-SCUplinkPortProfileSet for the port profile set created in VMM
$PortProfileSetId="fd9e4c9a-4ffa-4845-808d-930e6616b62f"

$vswitch=Get-VMSwitch -Name $VirtualSwitchName
$VMMPortFeatureId="1f59a509-a6ba-4aba-8504-b29d542d44bb"
$defaultPortFeature = Get-VMSwitchExtensionPortFeature -FeatureId $VMMPortFeatureId
$VMMFeatureId="8b54c928-eb03-4aff-8039-99171dd900ff"
$currentFeature = Get-VMSwitchExtensionSwitchFeature -SwitchName $VirtualSwitchName -FeatureId $VMMFeatureId
$defaultFeature = Get-VMSwitchExtensionSwitchFeature -FeatureId $VMMFeatureId
$defaultFeature.SettingData.LogicalSwitchId=$LogicalSwitchId
$defaultFeature.SettingData.LogicalSwitchName=$VirtualSwitchName
Add-VMSwitchExtensionSwitchFeature -SwitchName $VirtualSwitchName -VMSwitchExtensionFeature $defaultFeature

$defaultPortFeature = Get-VMSwitchExtensionPortFeature -FeatureId $VMMPortFeatureId
$defaultPortFeature.SettingData.PortProfileSetId=$PortProfileSetId
$defaultPortFeature.SettingData.PortProfileSetName=$PortProfileName
$defaultPortFeature.SettingData.NetCfgInstanceId="{ " + $vswitch.Id + "}"
Add-VMSwitchExtensionPortFeature -SwitchName $VirtualSwitchName -VMSwitchExtensionFeature $defaultPortFeature
-ExternalPort
```

After you run the script, refresh the host in VMM and verify if VMM recognizes the switch as logical switch.

How to create logical switches

This article describes how to create logical switches in the System Center - Virtual Machine Manager (VMM) fabric, convert a host virtual switch to a logical switch, and set up virtual switch extensions if you need them.

A logical switch brings virtual switch extensions, port profiles, and port classifications together so that you can configure each network adapter with the settings you need, and have consistent settings on network adapters across multiple hosts. You can team multiple network adapters by applying the same logical switch and uplink port profile to them.

Set up virtual switch extensions

You install switch extensions on the VMM server and then include them in a logical switch. There are a few types of switch extensions:

- **Monitoring extensions** can be used to monitor and report on network traffic, but they cannot modify packets.
- **Capturing extensions** can be used to inspect and sample traffic, but they cannot modify packets.
- **Filtering extensions** can be used to block, modify, or defragment packets. They can also block ports.
- **Forwarding extensions** can be used to direct traffic by defining destinations, and they can capture and filter traffic. To avoid conflicts, only one forwarding extension can be active on a logical switch.

You can set up a virtual switch extension manager (network manager) if you want to managed extensions using a vendor management console and the VMM console together.

Set up a virtual switch extension manager

1. Obtain the provider software from your vendor and install the provider on the VMM management server. If

you have a cluster install it on all nodes.

2. Click **Fabric > Home > Show > Fabric Resources > Networking > Switch Extension Managers**.
3. In **Add Virtual Switch Extension Manager Wizard > General** specify the manufacturer and type the connection string for example myextmanager1.contoso.com:443. The exact syntax is defined by the vendor. Specify the account you want to use to connect to the resource.
4. In **Host Groups** specify the host groups for which you want to use the extension manager.
5. In **Summary** review settings and click **Finish**. Check that the extension appears in the **Virtual Switch Extension Managers** pane.

Set up a logical switch

1. Make sure you have at least one uplink port profile before you begin.
2. Click **Fabric tab > Networking > Logical Switches > Create Logical Switch**.
3. In **Create Logical Switch Wizard > Getting Started**, review the information.
4. In **General**, specify a name and optional description.
5. In **Uplink Mode**, select:
 - **No Uplink Team** if you're not using teaming.
 - **Embedded Team** if you want to deploy the switch with SET-based teaming
 - **Team** if you want to use NIC teaming
6. In **Settings**, select the minimum bandwidth mode. If you've deployed Microsoft network controller, you can specify that it should manage the switch. If you enable this setting you won't be able to add extensions to the switch.
7. Enable SR-IOV if you need to. SR-IOV enables virtual machines to bypass the switch and directly address the physical network adapter. If you want to enable:
 - Make sure that you have SR-IOV support in the host hardware and firmware, the physical network adapter, and drivers in the management operating system and in the guest operating system.
 - Create a native port profile for virtual network adapters that is also SR-IOV enabled.
 - When you configure networking settings on the host (in the host property called Virtual switches), attach the native port profile for virtual network adapters to the virtual switch by using a port classification. You can use the SR-IOV port classification that is provided in VMM, or create your own port classification.
8. In **Extensions**, if you're using virtual switch extensions select them and arrange the order. extensions process network traffic through the switch in the order you specify. Note that only one forwarding extension can be enabled.
9. In **Virtual Port** add one or more port classifications and virtual network adapter port profiles. You can also create a port classification and set a default classification.
10. In **Uplink** add an uplink port profile, or [create a new one](#). When you add an uplink port profile, it is placed in a list of profiles that are available through that logical switch. However, when you apply the logical switch to a network adapter in a host, the uplink port profile is applied to that network adapter only if you select it from the list of available profiles.
11. In **Summary** review the settings and click **Finish**. Verify the switch appears in **Logical Switches**.

Convert virtual switch to logical switch

If a host in the VMM fabric has a standard virtual switch with or without SET, you can convert it to use as a

logical switch.

NOTE

- Before you can convert, you need a logical switch in place, with specific settings.
- You must be a member of the Administrator user role, or a member of the Delegated Administrator user role, where the management scope includes the host group in which the Hyper-V host is located.

Compare switch settings

1. Record if NIC Teaming (LBFO) or SET is being used on the host.
2. If you are using NIC teaming on the host, record teaming and load balancing settings by running the PowerShell commandlet *Get-NetLbfoTeam*.
3. In **Hyper-V Manager**, right-click the host > **Virtual Switch Manager**. Select the virtual switch and verify whether **Enable single-root I/O virtualization (SR-IOV)** is selected. Close Hyper-V Manager.
4. In the VMM console > **Fabric** > **Servers** > **All Hosts**, right-click the host > **Properties**.
5. In **Virtual Switches**, note the properties, including logical network, and minimum bandwidth mode.
6. In **Fabric** > **Networking** > **Logical Switches**, right-click the logical switch that you want to convert the host configuration to, and click **Properties**.
7. In **Logical Switches**, record the information:
 - In **General**, record the uplink mode, whether SR-IOV is enable, and minimum bandwidth mode.
 - In **Extensions**, note whether any forwarding extensions have been added to the logical switch.
 - In **Virtual port**, record the names of the port profiles that are listed. Be sure to note if one of them has SR-IOV in the name.
 - In **Uplinks**, record the network sites, whether uplink mode is teamed, the load balancing algorithm, and teaming mode.
8. In **Fabric** > **Networking**, click **Port Profiles**. For any relevant port profiles, click **Properties**. In **Offload Settings**, see if **Enable Single-root I/O virtualization** is checked.
9. Now compare the recorded information that you recorded for the logical switch and port profiles, with the virtual switch information.
10. Review the following table to see whether you can convert the host to use the logical switch.

ITEM	CONVERSION
SR-IOV	The SR-IOV setting (enabled or disabled) must be the same in the logical switch as it is in the virtual switch. If SR-IOV is enabled, it must be enabled in the logical switch itself, and in at least one virtual network adapter port profile within the logical switch.
Uplink mode Load balancing algorithm Teaming mode	The Uplink mode setting must match. If the uplink mode is Team , then the Load balancing algorithm and Teaming mode must also match.
Minimum bandwidth mode	Must match.

ITEM	CONVERSION
Network sites	The logical switch must be configured for the correct network sites (in the correct logical network) for this host.

11. If the settings in the logical switch don't match as described in the table, you need to find or create a logical switch that does match.

Convert a host to use a logical switch

NOTE

- The conversion will not interrupt network traffic.
- If any operation in the conversion fails, no settings will be changed, and the switch will not be converted.

1. In VMM, click **Fabric > Servers > All Hosts**. Right-click the host > **Properties**.
2. On the **Virtual Switches** tab, click **Convert to Logical Switch**.
3. Select the logical switch that you want to convert the host to. Then select the uplink port profile to use, and click **Convert**.
4. The **Jobs** dialog box might appear, depending on your settings. Make sure that the job has a status of **Completed**, and then close the dialog box.
5. To verify that the switch was converted, right-click the host, click **Properties**, and then click the **Virtual Switches** tab.

Next steps

[Apply network settings](#) on a host with a logical switch.

Set up MAC address pools in the VMM fabric

12/7/2018 • 2 minutes to read

This article provides information about System Center - Virtual Machine Manager (VMM) default MAC addresses, and describes how to create and manage a custom MAC address pool.

VMM uses static MAC address pools to automatically generate and assign MAC address to VMs. This article describes default MAC address pools in the VMM fabric and explains how to create custom pools.

Default MAC address pool settings:

MAC POOL NAME	ENVIRONMENT	DEFAULT RANGE
Default MAC address pool	Hyper-V	00:1D:D8:B7:1C:00 – 00:1D:D8:F4:1F:FF
Default VMware MAC address pool	ESX/ESXi	00:50:56:00:00:00 – 00:50:56:3F:FF:FF

Before you start

Before you create a custom MAC pool note that:

- If you want to divide one of the default pools into smaller custom pools, you must first delete the default MAC address pool or the default VMware MAC address pool. You must delete the default pool to avoid duplicate MAC address assignments.
- The first three octets of the beginning and ending MAC address must be the same.
- You must enter a valid hexadecimal values between 00 and FF.
- The ranges that you specify cannot overlap.
- The address range must not have the multi-cast bit set to 1. For example, you cannot use addresses that start with X1, X3, X5, X7, X9, XB, XD, or XF, where X is any value.
- To avoid conflicts with addresses reserved by Microsoft, VMware, and Citrix, do not use the following prefixes:
 - Reserved for Microsoft: 00:03:FF; 00:0D:3A; 00:12:5A; 00:15:5D; 00:17:FA; 00:50:F2; 00:1D:D8 (except for the 00:1D:D8:B7:1C:00 – 00:1D:D8:F4:1F:FF range that is reserved for VMM)
 - Reserved for VMware: 00:05:69; 00:0C:29; 00:1C:14; 00:50:56 (except for the 00:50:56:00:00:00 – 00:50:56:3F:FF:FF range that is the reserved as the default VMware static range)

Create a custom pool

1. Click **Fabric > Networking > MAC Address Pools > Home > Show > Fabric Resources > Create > Create MAC Pool**.
2. In **Create MAC Address Pool Wizard > Name and Host Group** specify a name and description. In **Host Group** select the host groups that should use the pool.
3. In **MAC Address Range** specify the start and end addresses.
4. In **Summary** review the settings and click **Finish**. When the job shows as **Completed** verify pool in **MAC Pools**.

Release IP addresses

In some circumstances you might want to remove addresses from the MAC pool. For example if a host that was assigned an IP address during bare metal deployment is removed from VMM management, or if a VM goes into a

missing state because it was removed outside VMM.

1. Click **Fabric > Networking > MAC Address Pools > Home > Show > Fabric Resources**.
2. In **MAC Pools** click the pool you want to modify > **Properties**.
3. In **Inactive addresses** select the addresses you want to release.

Next steps

Learn about [creating an IP address pool](#).

Integrate load balancing with VMM service templates

5/30/2019 • 8 minutes to read

Read this article to learn about integrating Windows network load balancing (NLB), and hardware load balancers with System Center - Virtual Machine Manager (VMM) service templates.

Service templates group VMs together to provide an app. They contain information about a service, including the VMs that are deployed as part of the service, the applications installed on VMs, and the network settings that should be used. You can add VM templates, network settings, applications, and storage to a service template.

Service templates can be single or multi tier. A single tier service contains one VM used as a specific app. A multi-tier service contains multiple VMs. [Learn more](#).

Set up load balancing for a service tier

You can add a load balancer to load balance requests to VMs in a service tier. You can use a hardware load balancer, or NLB for round robin balancing.

To add a load balancer, you'll need to do the following actions:

- Ensure you have [logical networks configured](#). The logical networks should have associated network sites. Those network sites should have one or more associated subnets from which you can create static IP address pools. In addition associated each network site with the host group where the service will be deployed.
- [Create an IP address pool](#) for the logical networks. The IP pool must contain a reserved range of virtual IP addresses that can be assigned to the load balancer. You must set up the static IP address pools for the load balancer and for the virtual machines behind the load balancer. These can be from the same pool or from different pools, but you'll need both VIPs and IP addresses for the virtual machines.
- [Create VM networks](#) on top of logical networks.
- [Create VIP templates](#): A virtual IP (VIP) template contains load balanced settings for a specific type of network traffic. After you create a VIP template, you can specify it when you set up load balancing in a service template.
- [Set up a hardware load balancer](#): If you want to enable hardware load balancing in a service template, there are a number of prerequisites you'll need to prepare.
- [Set up NLB](#): If you don't want to use a hardware load balancer, you can use NLB. There are some requirements and limitations.

Create VIP templates

1. In the VMM console click **Fabric > Networking > VIP Templates**.
2. Click **Home > Show > Fabric Resources > Create > Create VIP Template**.
3. In the **Load Balancer VIP Template Wizard > Name**, specify the template name and description. In **VIP port**, specify the port that will be used for the type of network traffic you want to balance. For example, 443 for HTTPS traffic. In **Backend port**, specify the port on which the backend server is listening for requests.
4. In **Type** do the following:
 - To use NLB click **Microsoft** in the manufacturer list and **Microsoft network controller** in **Model**.
 - To use a hardware load balancer click **Generic** to create a template for any supported hardware load balancer. Click **Specific** to create a template for a specific load balancer and specify the manufacturer and model.
5. In **Protocol**, click the protocol for which you want to create the VIP template.

- If you select **HTTPS**, you'll need to specify where the traffic terminates.
 - Select **HTTPS passthrough** to pass the traffic to the VM without decrypting it.
 - Select **HTTPS terminate** to terminate and decrypt the HTTPS traffic at the load balancer. This option gives the load balancer more information such as cookies and headers. To use this option specify the subject name of a certificate on the load balancer that can be used for HTTPS authentication. With this option, you can enable **Re-Encrypt** to reencrypt the HTTPS traffic from the load balancer to the VM.
 - Select **Custom** to specify **TCP**, **UDP**, or both.
6. In **Persistence**, select **Enable persistence** to make the client session sticky (affinity). This setting means that the load balancer will always try to direct the same client to the same VM. It's based on the specified source IP address and subnet mask, the destination IP address, and other parameters that vary depending on the protocol.
 7. In **Health Monitors**, you can optionally specify that a verification should run against the load balancer at regular intervals. To add a health monitor specify the protocol and the request. For example entering the command GET? makes an HTTP GET request for the home page of the load balancer and checks for a header response. You can also modify the response type, and monitoring interval, timeout, and retries. Note that the timeout should be less than the interval.
 8. In **Load Balancing**, select which load balancing method you want to use. You can configure new connections to be directed based on the least connections or the fastest response time, using round robin, or using a custom method supported by the load balancer. If you're enabling NLB, select **Round Robin**.
 9. On the **Summary** page, review the settings and click **Finish**. The **Jobs** dialog appears. Wait for a **Completed** status. Then verify that the template appears in the **VIP Templates** pane.

Set up a hardware load balancer

Set up a hardware load balancer as follows:

- **Get a configuration provider:** To add a supported hardware load balancer, you'll need to download and install a configuration provider available from the load balancer manufacturer. VMM currently supports [Brocade ServerIron ADX load balancer provider](#), and [Citrix NetScaler load balancer provider](#). The provider is a VMM plug-in that translates VMM PowerShell commands to the load balancer API. After you've installed the provider, you should restart the VMM service (**net stop scvmmsservice > net start scvmmsservice**).
- **Set up an account:** Create a VMM Run As account with a user name and password with permissions to configure the downloaded load balancer.
- **Add the load balancer to VMM:** You add a hardware load balancer to VMM using the Add Load Balancer Wizard.

Add the hardware load balancer to VMM

During the wizard, you select the host groups for which the load balancer is available, specify the load balancer model, specify the address and port used to manage the load balancer, specify affinity to VMM logical network, select the configuration provider, and test the connection. You'll need to configure the hardware load balancer before you deploy a service. After the service is deployed a load balancer can't be added.

1. Click **Fabric > Networking > Load Balancers > Fabric Resources > Home > Add > Add Resources > Load Balancer**.
2. In **Add Load Balancer Wizard > Credentials**, select the Run As account with the load balancer credentials.
3. In **Host Group**, select each host group where the service will be deployed. Hosts should be able to access the load balancer. In addition, a physical network adapter on the host should be configured to use the same logical network as the service tier.
4. In **Manufacturer and Model**, select the appropriate entries.

5. In **Address**, specify the **IP address** and **FQDN** or **NetBIOS** names of the load balancer. Specify the port on which the load balancer listens for requests.
6. In **Logical Network Affinity** specify the affinity to logical networks. Note that:
 - For frontend affinity, you'll select the logical network from which the load balancer obtains its VIP. The VIP is the IP address that's assigned to the load balancer when you deploy it in a service template.
 - For frontend affinity, based on the logical networks VMM determines the static IP address pools that are accessible from both the load balancer and from the relevant host group
 - When selecting logical networks for frontend affinity, the associated network site with the reserved VIP address range should be available to the host groups associated with the load balancer.
 - For backend affinity you'll select the logical networks to which you want to make the load balancer available for connections from the VMs in a service tier.
7. In **Provider**, click the load balancer provider. Click **Test** to check the configuration.
8. In **Summary**, verify the settings and click **Finish**. The **Job** dialog box appears. Wait for a **Completed** status and check in the **Provider** column that the provider is active.

Set up NLB

NLB is automatically included as a load balancer in VMM. As long as you've set up an NLB VIP template no other action is required, but note that:

- NLB can't be used if VM networks are configured with network virtualization.
- NLB can't be used in service tiers running Linux VMs.

Enable load balancing

1. If the service template isn't open click **Library > Templates > Service Templates** and open it.
2. Click **Actions > Open Designer**.
3. In the Service Template Designer, click the **Service Template Components** group > **Add Load Balancer**.
4. Click the load balancer object. You'll identify it with the VIP template name.
5. Click **Tool > Connector**. Click the **Server connection** associated with template and then click a **NIC** object to connect the load balancer to the adapter. In the NIC properties check the address types and that the MAC address is static.
6. With the **Connector** enabled click the **Client connection** associated with the load balance and then click a logical network object.
7. Save the service template in **Service Template > Save and Validate**.

Set up the hardware VIP for user access

When the service is deployed, VMM automatically selects a VIP from the reserved range in the static IP address pool, and assigns it to the load-balanced service tier. To enable users to connect to the service, after the service is deployed you need to determine the VIP and configure a DNS entry for it.

1. After the service is deployed click **Fabric > Networking > Load Balancers**.
2. Click **Show > Service > Load Balancer Information for Services** and expand the service to see which VIP is assigned.
3. Request that the DNS administrator manually create a DNS entry for the VIP. The entry should be the name that users will specify to connect to the service. For example, servicename.contosol.com.

Next steps

You can also [set up a software load balancer](#) in a software defined networking (SDN) infrastructure in the VMM fabric.

Set up an IPAM server in the VMM fabric

This article explains how to add an IP Address Management (IPAM) server to the System Center - Virtual Machine Manager (VMM) networking fabric.

An IPAM server helps you to plan, track, and manage the IP address space used in your networks.

- With an IPAM server in the VMM fabric, the IP address settings that are associated with logical networks and VM networks in VMM are synchronized using settings stored on the IPAM server.
- As an administrator you can use the IPAM server to configure and monitor logical networks and their associated network sites and IP address pools. You can also use the IPAM server to monitor the usage of VM networks that you have configured or changed in VMM.
- Tenants must continue to use the VMM server (not IPAM) to configure VM networks that use network virtualization. In other words, to control the address space that is typically controlled by tenants rather than by VMM administrators.

Before you start

- Make sure you have an IPAM server. [Learn more](#). The IPAM server can be running these versions of [Windows Servers](#).
- Create or identify a domain account and set it to never expire. On the IPAM server add the account to these groups:
 - **IPAM ASM Administrators**: A local group that exists on all IPAM servers, and provides permissions for IP address space management (ASM). For more information, see [Assign Administrator Roles](#).
 - **Remote Management Users**: A built-in group that provides access to WMI resources through management protocols, such as WS-Management through the Windows Remote Management service.
- Check that the time is synchronized on the IPAM and VMM servers. This depends on settings for the Windows Time Service. If you can't synchronize them you'll need to update permissions on the IPAM software so that VMM can query the current time setting on the server. To do this, on the IPAM server run **mimgmt.msc** to open the WMI Control (Local) snap-in. Right-click **WMI Control (Local)** > **Properties** > **Security**. Navigate to **Root\CIMV2**, click the Security button Security and select the account you configured. For **Remote Enable**, select **Allow**.
- Verify the FQDN of the IPAM server to use as a connection string.
- Verify the names of the VMM host groups for which you want to use the IPAM server.
- The provider software for an IPAM server is included in VMM. You don't need to install it. You can review settings in **Settings > Configuration Providers**.
- If you want to use the IPAM server to delete a logical network, delete the IP address subnets assigned to that logical network, and do not delete the name associated with the **VMM Logical Network** field on the IPAM server. The two servers will then be able to synchronize correctly, and the logical network will be deleted. If you do delete the name associated with the **VMM Logical Network** field on the IPAM server, you must go to the VMM server and delete the network sites and the logical network. Then, after the two servers synchronize, the deletion will be complete.

Add an IPAM server to the fabric

1. Click **Fabric > Home > Show > Fabric Resources > Fabric > Networking > Network Service**. Network services include gateways, virtual switch extensions, network managers (which include IPAM servers), and top-of-rack (TOR) switches.
2. Click **Home > Add > Add Resources > Network Service**.
3. In **Add Network Service Wizard > Name** specify a name and optional description.
4. In **Manufacturer and Model > Manufacturer** click **Microsoft**, and click **Model > Microsoft Windows Server IP Address Management**.
5. In **Credentials** page specify the account you created.
6. In **Connection String** page, in the **Connection string** box, type the FQDN of the IPAM server. If you've configured a specific port on the IPAM server, end the string with the port number (for example, **:443**). If a port number is not specified, the default port for the IPAM server is used.
7. In **Provider > Configuration provider > Microsoft IP Address Management Provider**, click **Test** to run basic validation tests with the provider. Results that say **Passed** or **Failed** indicate whether the provider works as expected. One possible cause of failure is insufficient permissions in the Run As account. Results that say **Implemented** and **Not implemented** are informational only, and indicate whether the provider supports a particular API.
8. In **Host Group** select one or more host groups for which you want integration between the IPAM server and the VMM server.
9. In **Summary** review the settings and click **Finish**. Check that the IPAM server is listed under **Network Services**. Right-click the server > **Refresh** to get the latest settings.
10. On the IPAM server, to view the logical networks and related settings that were configured in VMM, navigate to **VIRTUALIZED IP ADDRESS SPACE**, and then to **Provider IP Address Space**. For each logical network, the IPAM server will have an address space (an overarching category that is found in IPAM, but not in VMM) with a name that is based on the name of the logical network. The logical network will be contained within the address space, with the name of the logical network displayed under the heading **VMM Logical Network**. To see the types of information that are stored in IPAM, expand the address space and select different views.

The following table can help you interpret some of the information that you see on the IPAM server:

VMM NAME	IPAM NAME
Logical network	VIRTUALIZED IP ADDRESS SPACE Provider IP Address Space: VMM Logical Network column
Network site	VIRTUALIZED IP ADDRESS SPACE Provider IP Address Space: Network Site column
IP address subnet	IP Address Subnet (same name in IPAM as in VMM)
IP address pool	IP Address Range
VM network	VIRTUALIZED IP ADDRESS SPACE Customer IP Address Space: VM Network column

Next steps

[Set up logical networks](#)

Set up the VMM storage fabric

5/6/2019 • 6 minutes to read

You can use System Center - Virtual Machine Manager (VMM) to manage your physical and virtualized infrastructure. As part of that management VMM can manage storage that you assign to virtual hosts and clusters, and VMs.

- **Local and remote:** VMM recognizes both local and remote storage. Local is storage located on the VMM server or directly attached to it, commonly a disk drive on the server that's connected with inbuilt RAID or SAS JBOD connectivity. Obviously this type of dedicated host storage isn't shared, and doesn't provide resilience or high availability.
- **Block and file-based:** VMM can manage block storage devices and file-based storage.

Block storage

FEATURE	DETAILS
Connection	<p>VMM can manage block storage devices that connect using Fibre Channel, Serial Attached SCSI (SAS), or iSCSI (Internet SCSI).</p> <p>VMM can discover and manage iSCSI arrays with static, dynamic, or manual targets.</p> <p>You can use a Microsoft iSCSI target server as a storage device by installing its provider.</p>
Protocols	<p>VMM provides support for storage devices that use the SMI-S and SMP protocols.</p> <p>VMM uses Window Storage Management API (SMAPI) to manage block-based storage devices compliant with SMI-S or SMP specifications.</p> <p>VMM combines SMAPI and SMP to manage directly attached storage and external storage arrays.</p> <p>VMM combines SMAPI and the Storage Management service (which functions as an SMI-S client) to manage SMI-S storage devices.</p> <p>Vendors of storage devices with the SMI-S standard create SMI-S providers for their devices.</p>
Virtualization hosts	<p>Storage configured in VMM can only be used for Hyper-V hosts and clusters.</p>
Supported storage arrays	<p>Specific storage arrays are supported in VMM. You can use other arrays, but there's no guarantee that you'll be able to manage all storage tasks in VMM. We recommend that you chat to your storage provider to determine VMM support.</p>

FEATURE	DETAILS
Virtual fibre channel	<p>If you want to use virtual fibre channel to give VMs direct access to fibre channel storage, you can manage this storage with VMM in these configurations.</p> <p>Single storage array connected to a single fabric (comprised of single or multiple switches) connected to a single vSAN.</p> <p>Single storage array connected to multiple fabrics (comprised of single or multiple switches per fabric) connected to a single vSAN.</p> <p>Multiple storage arrays connected to a single fabric (comprised of single or multiple switches) connected to a single vSAN.</p> <p>Multiple storage arrays connected to multiple fabrics (comprised of single or multiple switches per fabric) connected to multiple vSANs. This configuration provides dual-redundant paths to storage arrays.</p> <p>A vSAN can only include HBAs from a single fabric.</p>

Set up block storage

The general process for setting up block-based storage in the VMM fabric is as follows:

1. **Create storage classifications:** You create storage classifications to group storage based on shared characteristics, often performance and availability. Then instead of assigning specific storage devices to them, you assign storage to VMM host groups you assign a specific classification so that host groups can use any available storage device with the assigned classification. You don't have to create classifications before you add storage devices. You can create them during storage device discovery.
2. **Add the storage:** You add the storage as a resource in the VMM fabric. When you add the device VMM automatically discovers any existing storage pools and logical units on the device. You can classify storage as you add it.
3. **Configure storage and allocate capacity:** After a storage array is managed by VMM you can configure settings. You can specify how you want to use rapid provisioning on the device (snapshots or cloning). You can add and modify storage pools and storage logical units (LUNs) in the pools. You can allocate capacity (either entire storage pools) or specific LUNs to one or more host groups.
4. **Use the storage:** After storage is allocated to a host group you can use the storage for a specific host or cluster. When you add a host or cluster to a host group the host and cluster can use storage associated with the group.

File storage

VMM can manage file storage that supports the SMB 3.0 protocol. SMB is supported by file shares on computers running Windows Server 2012 or higher, and by third-party vendors of network-attached storage (NAS) devices.

- **Windows file server:** You can add a remote file server as a storage device or you can scale file-based storage Scale-Out File Server (SOFS).
- **Scaled-out file server (SOFS):** SOFS provides a file server cluster in which storage is shared between the cluster nodes. Storage for SOFS could be a SAN (SAS, iSCSI, Fibre Channel) or could integrate with Storage Spaces Direct.
- **Storage Spaces Direct (S2D):** S2D is the next evolution of Microsoft Storage Spaces, which virtualizes storage by grouping disks into storage pools and creating virtual disks (storage spaces) from the pool capacity. In S2D you can build highly available storage using local storage. This removes the need for remote SAN

storage devices and enables use the storage devices that weren't previously available, such as SATA SSD or NVMe flash. [Learn more.](#)

- **Storage replication:** VMM supports Windows Storage Replica for protecting data in a primary storage volume that replicating it to a secondary volume. [Learn more.](#)
- **Storage resources:** You can control access to shared storage on an SOFS or VM by setting storage quality-of-service (QoS) policies. These policies set maximum and minimum bandwidth for storage resources.

Set up file storage

The general process for setting up file storage in the VMM fabric is as follows:

1. **Add and discover storage:** You add the file server as a resource in the VMM fabric. When you add the device VMM automatically discovers any file shares on the device. You can classify storage as you add it.
2. **Create storage classifications:** You create storage classifications to group file shares based on shared characteristics, often performance and availability. Then instead of assigning specific storage devices to VMM host groups you assign a specific classification so that host groups can use any available storage with the assigned classification. You don't have to create classifications before you add storage devices. You can create them during storage device discovery.
3. **Provision storage:** After a file server is managed by VMM you can configure settings. For example you can modify storage pools on an SOFS or create a file share.
4. **Allocate capacity:** After you have the storage set up you can allocate capacity from the array. You allocate capacity by assigning file shares to one or more host groups.
5. **Use the storage:** After storage is allocated to a host group you can use the storage for a specific host or cluster. When you add a host or cluster to a host group the host and cluster can use storage associated with the group.
6. **Decommission storage:** VMM can decommission the storage it manages.

[Learn more](#) about setting up file storage in VMM.

Storage classifications

Storage classifications provide a layer of abstraction over specific storage devices. You group storage devices together based on their characteristics. For example you could create:

- Bldg1Gold: A set of solid-state drives (SSDs) that you will make available to users in building 1.
- Bldg1Silver: A set of SSDs and hard disk drives (HDDs) that you will make available to users in building 1.
- Bldg2Gold: A set of SSDs that you will make available to users in building 2.
- Bldg2Silver: A set of SSDs and HDDs that you will make available to users in building 2.

After you've created classifications you assign them to storage pools that include block or file-based storage. You can tweak classification settings for file shares within pools as required.

Monitor storage health

This functionality helps you to monitor the health and operational status of a storage pool, LUNs and physical disks in the VMM fabric.

You can monitor storage health from VMM console, **Fabric** page.

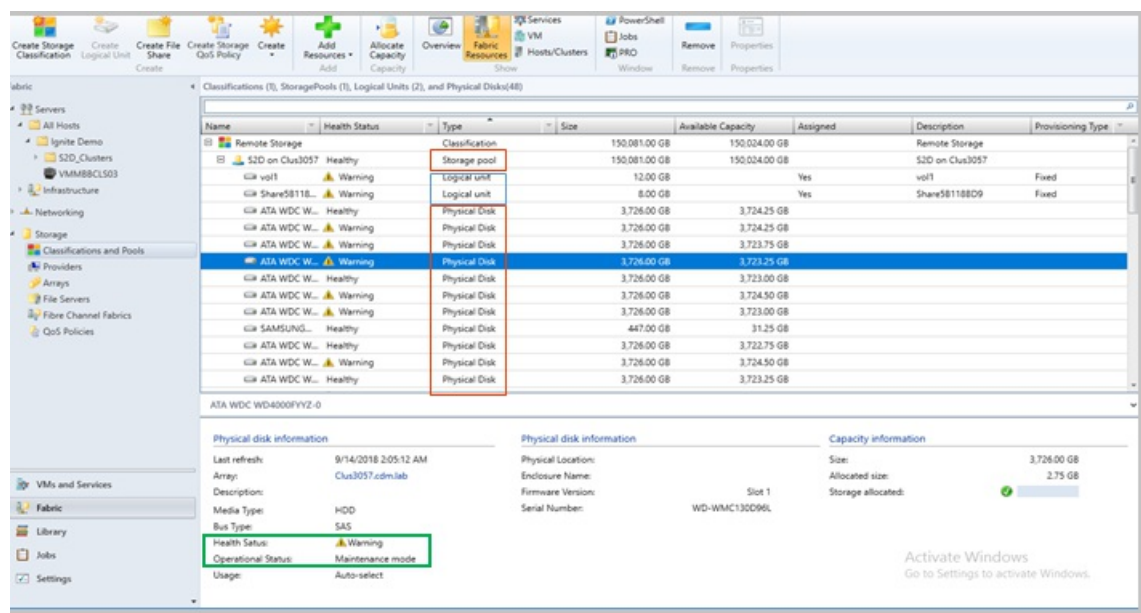
Follow these steps:

1. Go to VMM console, **Fabric > Storage > Classification and Pools**.

The **Health Status** column displays the status of the storage pool, LUN and Physical disks.

2. To check the operational status, select a storage item.

The information pane for the selected storage item displays the operational status as shown in the sample below:



Next steps

Set up storage classifications in the VMM fabric

Set up storage classifications in the VMM fabric

12/7/2018 • 2 minutes to read

Use storage classifications to abstract storage devices in the System Center - Virtual Machine Manager (VMM) fabric. You classify storage devices with similar characteristics, and assign these classifications, rather than specific storage devices, to hosts and clusters. The host and cluster can then use any available storage in the classification.

Classifications are often based on storage types or performance characteristics. For example, you could create:

NAME	DESCRIPTION
GOLD	Storage pool based on solid-state drives (SSDs) that delivers high performance for I/O intensive applications
SILVER	Fibre Channel Serial Attached SCSI (SAS) storage (RAID 5)
BRONZE	iSCSI Serial ATA (SATA) storage (RAID 5)

Create classifications:

1. Click **Fabric** > **Storage**, right-click **Classification and Pools** > **Create Classification**.
2. In **New Classification**, enter a name and description > **Add**.

Next steps

[Add storage devices](#) to the VMM fabric.

Add storage devices to the VMM fabric

12/7/2018 • 2 minutes to read

To manage storage in System Center - Virtual Machine Manager (VMM), you discover and add it to the VMM storage fabric.

Before you start

Make sure that the storage device is supported before you add it.

Add a storage device

1. Click **Fabric > Storage > Add Resources > Storage Devices**.
2. In **Add Storage Devices Wizard > Select Provider Type**, select to add a storage device with SMI-S or SMP, according to the device you're using.
3. In **Specify Discovery Scope**:
 - If you're using SMI-S, specify whether the provider uses **SMI-S CIMXML** or **SMI-S WMI WMI**, add the IP address/FQDN, and add the port used to connect to the provider on the remote server. You can enable SSL if you're using CIMXML. Then specify an account for connecting to the provider.
 - If you're using SMP, select the provider from the list. If it isn't in the list click **Import** to refresh it.
4. In **Gather Information**, VMM automatically tries to discover and import the storage device information. To retry click **Scan Provider**.
5. If you selected the option to use an SSL connection for an SMI-S provider, note that:
 - During discovery, the **Import Certificate** dialog box appears. Check settings and click **Import**. By default the certificate common name (CN) will be verified. This might cause storage discovery to fail if there's no CN, or it doesn't match.
 - If discovery fails because of the CN, disable CN verification in the registry on the VMM server. In the registry go to **HKEY_LOCAL_MACHINE/SOFTWARE/Microsoft/Storage Management/** and create a new DWORD value - **DisableHttpsCommonNameCheck**. Set the value to 1.
6. If the discovery process succeeds, the discovered storage arrays, storage pools, manufacturer, model, and capacity are listed on the page. When the process finishes, click **Next**.
7. In **Select Storage Devices**, you can specify a classification for each storage pool. Storage classifications group storage pools with similar characteristics together so that you can assign a classification as storage for a host or cluster, rather than a specific storage device. Learn more about setting up classifications.
8. On the **Summary** page, confirm the settings, and then click **Finish**. The **Jobs** dialog box appears. When status is **Completed** you can verify the storage in **Fabric > Storage**.

Next steps

[Allocate storage to host groups](#)

Allocate storage to host groups

12/7/2018 • 2 minutes to read

After block storage has been discovered and classified in the System Center - Virtual Machine Manager (VMM) fabric, you can allocate it to host groups. You can allocate an entire storage pool, or a specific logical unit (LUN).

- **Allocate storage pools:** You can optionally allocate storage pools to host groups. If you do you can:
 - Create and assign LUNs directly from Hyper-V hosts in the host group.
 - Use the storage pool for rapid provisioning using SAN cloning or snapshots. During this process you don't need to create LUNs because VMM request a copy of an existing LUN during provisioning.
- **Allocate LUNs:**
 - You can assign LUNs for storage pools that are managed in the VMM fabric.
 - In order to allocate LUNs to host groups there must be unassigned LUNs in the managed storage pools.
 - If you need additional LUNs you can create them outside of VMM in your storage management tool, or you can provision them directly in VMM if the storage pool is allocated to a host group.

Allocate a storage pool to a host group

1. Click **Fabric > Storage > Allocate Capacity**, and click the host group. If you're a delegate admin with scope restricted to host groups, right-click the host group > **Properties > Storage**.
2. The total and available storage capacity information is displayed for the host group. The storage capacity information includes the total and available capacity for both local and remote storage, and total and available allocated storage. Click **Allocate Storage Pools**.
3. Click a storage pool > **Add**.

Create a LUN in VMM

1. Make sure you've allocated the storage pool to the host group, and click **Fabric > Storage > Create Logical Unit**.
2. Specify the storage pool, a name an description for the LUN, and the size. Click **OK** to create the LUN.
3. Verify that the LUN was created in **Fabric Resources > Classifications, Storage Pools, and Logical Units**.

Allocate a LUN to a host group

1. Click **Fabric > Storage > Allocate Capacity > Allocate Storage Capacity** and click the host group.
2. Click **Allocate Logical Units**, select a unit > **Add**.

After LUNs are allocated to host groups you can assign them to Hyper-V hosts and clusters.

Next steps

After you set up Hyper-V hosts and clusters, learn about [provisioning VMs](#).

Set up a Microsoft iSCSI Target Server in the VMM storage fabric

12/7/2018 • 2 minutes to read

Microsoft iSCSI Target Server is a server role that enables a Windows server machine to function as a storage device. This article explains how to set up a Microsoft iSCSI Target Server in System Center - Virtual Machine Manager (VMM) storage.

Here's what you need to do:

1. **Install the role:** Install the iSCSI Target Server role (**Server Roles** > **File and Storage Services** > **File and iSCSI Services**) on a server that you want to add as a block storage device.
2. **Set up virtual iSCSI disks:** After installing the role you'll need to set up virtual iSCSI disks and connect to the servers you want. [Learn more](#).
3. **Install the provider:** If the iSCSI Target Server runs Windows Server 2012, you must install the SMI-S provider on it. The provider is located with the setup files in \amd64\Setup\msi\iSCSITargetSMISProvider.msi, and on the VMM server in \Program Files\Microsoft System Center 2012\Virtual Machine Manager\Setup\Msi\iSCSITargetProv\iSCSITargetSMISProvider.msi. You'll need to run the .msi file on the iSCSI Target Server. If the server's running Windows Server 2012 R2 you don't need to install the provider.
4. **Add account:** Add the VMM admin account as an administrator on the server.
5. **Discover in VMM:** [Add the storage device](#) to VMM. Select **SAN and NAS devices discovered and managed by a SMI-S provider** as the provider type, and specify the IP address or FQDN as the server. Select the account with permissions to the server as the Run As account. Add it to the required storage classification, and complete the **Add Storage Devices Wizard**.

After adding the server as a storage device under VMM management you can allocate the storage pools and LUNs to a host group and provision storage to hosts and clusters.

PowerShell example

You can use VMM to configure the iSCSI Target Server through Windows PowerShell. This section lists some common tasks with examples of Windows PowerShell commands that you can use for those tasks. The SMI-S provider supports all management tasks through VMM.

Manage storage on an iSCSI target server

Open PowerShell and use the cmdlets described below to manage iSCSI target server resources in VMM.

Add a storage provider

COMMAND	PURPOSE
<pre>\$Cred = Get-Credential</pre>	Obtain the iSCSI Target Server local administrative credentials that are based on user name and password. Note that any account that is part of the Local Administrators group is sufficient.
<pre>\$Runas = New-SCRunAsAccount -Name "iSCSIRunas" -Credential \$Cred</pre>	Create a Run As account in VMM.

COMMAND	PURPOSE
<pre>Add-SCStorageProvider -Name "Microsoft iSCSI Target Provider" -RunAsAccount \$Runas -ComputerName "<computername>" -AddSmisWmiProvider</pre>	Add the storage provider.

View storage properties

COMMAND	PURPOSE
<pre>\$array = Get-SCStorageArray -Name "<computername>"</pre>	Review the storage array attributes.
<pre>\$array.StoragePools</pre>	View available storage pools.

Add pools from iSCSI Target Server to VMM management

COMMAND	PURPOSE
<pre>\$pool = Get-SCStoragePool -Name "MS iSCSITarget Concrete: D:"</pre>	Get the specific storage pool to add.
<pre>\$class = New-SCStorageClassification -Name "gold"</pre>	Create a storage classification, if none exists.
<pre>Set-SCStorageArray -AddStoragePoolToManagement \$pool -StorageArray \$pool.StorageArray -StorageClassification \$class</pre>	Add the storage pool to VMM.
<pre>Set-SCStoragePool -StoragePool \$pool -AddVMHostGroup (Get-SCVMHostGroup -Name "All Hosts")</pre>	Allocate the storage pool to a virtualization server group.

Create a LUN

COMMAND	PURPOSE
<pre>\$LUN = New-SCStorageLogicalUnit -Name "iSCSI1" -StoragePool \$pool -DiskSizeMB 1000</pre>	Create an iSCSI logical unit number (LUN).
<pre>Set-SCStorageLogicalUnit -StorageLogicalUnit \$LUN -VMHostGroup (Get-SCVMHostGroup -Name "All Hosts")</pre>	Allocate the LUN to the host group.
<pre>\$host = Get-SCVMHost -ComputerName <host name></pre>	Retrieve the properties of a host.
<pre>Register-SCStorageLogicalUnit -StorageLogicalUnit \$LUN -VMHost \$host</pre>	Assign the LUN to the host.

Decommission resources

COMMAND	PURPOSE
<pre>Remove-SCStorageLogicalUnit -StorageLogicalUnit \$LUN</pre>	Delete a LUN.
<pre>Remove-SCStorageProvider -StorageProvider (Get-SCStorageProvider -Name "Microsoft iSCSI Target Provider")</pre>	Remove a storage provider.

Next Steps

Learn about [provisioning storage](#) for Hyper-V hosts and clusters.

Set up Hyper-V virtual fibre channel in the VMM storage fabric

12/7/2018 • 9 minutes to read

Read this article to set up Hyper-V virtual fibre channel in the System Center - Virtual Machine Manager (VMM) storage fabric.

Virtual fibre channel provides Hyper-V VMs with direct connectivity to fibre channel-based storage. Hyper-V provides fibre channel ports within guest operating systems, so that you can virtualize applications and workloads that have dependencies on fibre channel storage. You can also cluster guest operating systems over fibre channel.

Before you start

- VMM supports the following virtual fibre channel deployments:
 - Single storage array connected to a single fabric (comprised of single or multiple switches) connected to a single virtual SAN (vSAN). A vSAN is a named group of physical fibre channel Host Bus Adapter (HBA) ports on a host computer that VMs connect to access fibre channel storage devices.
 - Single storage array connected to multiple fabrics (comprised of single or multiple switches per fabric) that are connected to a single vSAN.
 - Multiple storage arrays connected to a single fabric (comprised of single or multiple switches) that is connected to a single vSAN.
 - Multiple storage arrays connected to multiple fabrics (comprised of single or multiple switches per fabric) that are connected to multiple vSANs. This configuration provides dual-redundant paths to storage arrays.

Here's what you need:

- One or more vSANs can be created for each host computer. A vSAN can only contain HBAs from a single fabric.
- Storage arrays, switches, and HBAs must have the latest firmware and drivers installed.
- Make sure that storage arrays can present logical units (LUs).
- Enable NPIV on fibre channel switches and HBAs.
- Hyper-V hosts must be running at least Windows Server 2012.
- Ensure that an SMI-S provider is installed. VMM manages fibre channel fabrics and SAN devices using the SMI-S provider. Remember not to install the SMI-S provider on the VMM server, but on a server that the VMM server can connect to with an FQDN or IP address.

Deploy virtual fibre channel

Here's what you need to do:

1. Discover and classify fibre channel fabrics.
2. Create vSANs for each host computer by grouping host HBA ports.
3. Create a VM that can access the virtual fibre channel storage.
4. Create zones that connect each host or VM vHBA to a storage array. Zones are used to connect a fibre channel array to a host computer VM.
5. Create LUNs and register them for a host, VM, or service tier.
6. Create a service template, add VM templates to it. For each vHBA, specify dynamic or static WWN assignments

and select the classification. Create and deploy a service tier based on the service template, to access Virtual Fibre Channel storage. You zone a fibre channel array to the tier, add a disk, create a LUN, and register the LUN to the tier.

Discover and classify fibre channel fabrics

1. Click **Fabric > Storage > Add Resources > Storage Devices**.
2. In **Add Storage Devices Wizard > Select Provider Type** select **Fibre Channel fabric discovered and managed by an SMI-S provider**.
3. In **Specify Discovery Scope** specify the IP address or FQDN, and the port number of the provider.
4. If you're using SMI-S specify whether the provider uses **SMI-S CIMXML** or **SMI-S WMI** and add the IP address/FQDN and port used to connect to the provider on the remote server. If you're using CIMXML, you can enable SSL.
5. Specify an account for connecting to the provider.
6. In **Gather Information**, VMM automatically discovers and imports the fibre channel fabric information. If the discovery process succeeds, the discovered fabric name, switches, and fabric World Wide Node Names (WWNN) are listed on the page. When the process successfully completes, click **Next**. To retry the discovery process for an unsuccessful attempt, click **Scan Provider**.
7. If you selected the option to use an SSL connection for an SMI-S provider note that:
 - During discovery, the **Import Certificate** dialog box appears. Check settings and click **Import**. By default, the certificate's common name (CN) is verified. Storage discovery may fail if there's no CN or it doesn't match.
 - If discovery fails because the CN disable CN verification in the registry on the VMM server. In the registry, go to **HKEY_LOCAL_MACHINE/SOFTWARE/Microsoft/Storage Management/** and create a new DWORD value - **DisableHttpsCommonNameCheck**. Set the value to 1.
8. On the **Fibre Channel Fabrics** page, do the following for each storage fabric that requires a classification:
 - a. In the **Storage Device** column, select the check box next to a fibre channel fabric that you want VMM to manage.
 - b. In the **Classification** column, select the classification that you want to assign to the fabric. The fabric classification task is separate from that for storage classification, although the concept is similar.
9. On the **Summary** page, confirm the settings, and then click **Finish**.

Create vSANs and assign HBAs

You can create vSANs and assign HBAs to it. One or more vSANs can be created for each host computer. Each vSAN can only contain HBAs that are from the same fabric.

Virtual Host Bus Adapters (vHBAs) represent the virtualization of fibre channel HBAs, and are used by VMs to connect with vSANs. Each vHBA has a World Wide Node Name (WWNN), which is different than the host HBA WWNN. Using N_Port ID Virtualization (NPIV), a host computer HBA can map to multiple vHBAs. HBA ports assigned to a vSAN can be added or removed as needed.

1. Click **Fabric** right-click the applicable host > **Properties > Hardware > New Virtual SAN**.
2. In **New Virtual SAN** specify a name and optional describes. In **Fibre Channel adapters**, select the check boxes next to the fibre channel adapters (HBAs) that you want to assign to the vSAN. Click **OK**.
3. If you want to edit vSAN port assignments, in **Properties > Hardware > FC Virtual SAN > Fibre Channel adapter details** select or unselect the HBA ports.
4. If you want to add a new vHBA and assign it to a vSAN, click **Properties > Hardware Configuration > New > Fibre Channel Adapter**. In **Virtual SAN name** select a vSAN to assign it. Specify whether you want to assign port settings for the vHBA statically or dynamically.
5. If you want to change the global default port settings for vHBA click **Properties > Hardware > Global**

Settings and modify settings in **Fibre Channel adapter details**. Note that changing these settings does not affect vHBA ports that have already been created. To apply a new setting to an existing vHBA port, recreate the port by removing it and then adding it again.

Create a VM template

vHBAs are used by VMs to connect with vSANs. In order for vHBAs to connect to vSANs, they first must be added to the hardware profile of a VM template.

1. Use the **Create Virtual Machine Wizard** to create a new VM, and then add a new fibre channel adapter (vHBA) to the **Configure Hardware** page of the VM template. For each vHBA that you create, specify dynamic or static WWPN assignments and select the fabric classification.
2. Still using the **Create Virtual Machine Wizard**, place and deploy the VM to a destination host. Make sure the host contains a virtual SAN that matches the storage fabric.

After you deploy the VM to a host, you can zone a virtual fibre channel storage array to the VM. Then, you create a LUN and register (unmask) it to the VM.

Create zones

Zones are used to connect a fibre channel array to a host or virtual machine (VM). The storage array target ports are mapped to the HBA ports on the host or to the virtual HBA (vHBA) ports for the VM. You can create zones for a host, a VM, or both. For Hyper-V failover clusters, a zone is needed for each host in the cluster. Note that:

- Zones are grouped into zonesets, which use common fibre channel fabric devices. When all zones in a zone set have been added, modified, or removed as needed, the zoneset must be activated. Zoneset activation pushes information for each zone down to the fibre channel switches in the selected fabric.
- Only members of the same zone can communicate with each other.
- You'll need to create new zones and then activate the zoneset. Activating a zoneset may cause some downtime in the fabric as information is propagated to all the switches.
- If you want to add a storage array to a Hyper-V cluster, you need to zone the array to each host computer first. Similarly, if you want to add an array to a guest cluster, you need to zone the array to each VM first.

Set up zones as follows:

1. Click **VMs and Services > Services**, right-click the applicable VM, > **Properties > Storage > Add > Add Fibre Channel Array**.
2. In **Add Fibre Channel Array > Properties** page > **Create New Zone** specify a zone name, select a storage array, and in **Fabric** select a switch. In **Storage array target ports**, select the applicable WWPN port or ports. In **Virtual machine initiator**, select the applicable WWPN port or ports. Then click **Create**. Click **Show aliases** to view the available zone aliases.
3. To activate the zoneset, click **Fabric > Name**, select the inactive zoneset > **Activate Zoneset**.
4. You can view the zonesets for a fabric in **Fabric > Fibre Channel Fabric > Name**, right-click the applicable fabric > **Properties > Zonesets**.
5. If you want to modify zoning for a storage array click **VMs and Services > applicable host > Properties > Storage > Fibre Channel Arrays > Edit > applicable array** and modify the zoning settings.

Create and register LUNs

For a host computer, VM, or computer service tier to access storage array resources, LUNs must be created and then registered (unmasked) to the host, VM, or tier.

1. Click **Fabric > Storage > Classifications and Pools**. Under **Name**, click the applicable storage device > **Create Logical Unit**.

2. In the **Create Logical Unit** select a storage pool, specify a name and description, and LUN size. Specify whether you want to create a thin or fixed size LUN.
3. To register the LUN, in **VMs and Services** pane, right-click the applicable VM > **Properties** > **Add** > **Add Disk**.
4. In **Create Logical Unit** select a storage pool, name, and size. Click **OK** to register the LUN.

Create and deploy a service tier

1. Using the **Service Template Designer**, create a service template and add the applicable VM templates you previously created to the service template.
2. Add a new virtual fibre channel adapter (vHBA) to the **Configure Hardware** page of the service template. For each vHBA that you create, specify dynamic or static WWPN port assignments and select the fabric classification.
3. Create service tier from the service template and assign the service tier to a computer tier.
4. Deploy the tier.
5. After you deploy it, you can zone a virtual fibre channel storage array to the service tier. Then create a LUN for the array and register (unmask) it to the tier.

Next steps

[Set up storage](#) for Hyper-V hosts and clusters.

Set up file storage in the VMM fabric

12/7/2018 • 3 minutes to read

You can manage file storage that supports SMB 3.0 in the System Center - Virtual Machine Manager (VMM) storage fabric. This includes Windows file servers, scale-out file servers (SOFS), network-attached storage (NAS) devices from third-parties such as EMC, and NetApp.

This article describes how to add file storage to the VMM fabric. After file shares are available in the fabric, you can assign them to Hyper-V hosts and clusters.

Add a file share to the VMM fabric

When you add a file server, VMM automatically discovers all the shares that are currently present on the server.

1. Click **Fabric > Storage > Add Resources > Storage Devices**.
2. In **Add Storage Devices Wizard > Select Provider Type**, select **Add a Windows-based file server as managed storage device** to manage a single or clustered file server in the VMM console.
3. In **Specify Discovery Scope**, specify the address or name of the file server. If the file server resides in a domain that isn't trusted by the domain in which the virtual machine hosts are located, select **This computer is in an untrusted Active Directory domain**. Select a Run As account that can access the file server.
4. In **Gather Information**, VMM automatically tries to discover and import information about the shares on the file server. If the discovery process succeeds, information about the file server is displayed. When the process finishes, click **Next**.
5. In **Select Storage Devices**, select the file shares that you want VMM to manage.
6. In **Summary**, review the settings and click **Finish**.

Create a file share

You need to create a file share on the file server. When you do this VMM assigns permissions automatically.

1. Click **Fabric > Storage > Providers**.
2. Select the file server > **Create File Share**.
3. In **Create File Share**, specify the path on which you want to create the share. If it doesn't exist VMM will create it.

Assign files shares

You can assign file shares on any host on which you want to create VMs that will use the file share as storage.

1. Click **Fabric > Servers > All Hosts**, and select the host or cluster node you want to configure.
2. In **Host > Properties**, click **Properties > Host Access**. Specify a Run As account. By default the Run As account that was used to add the host to VMM is listed. In the Run As account box, configure the account settings. You can't use the account that you use for the VMM service.
 - If you used a domain account for the VMM service account, add the domain account to the local Administrators group on the file server.
 - If you used the local system account for the VMM service account, add the computer account for the VMM management server to the local Administrators group on the file server. For example, for a VMM management server that is named VMMServer01, add the computer account VMMServer01\$.
 - Any host or host cluster that accesses the SMB 3.0 file share must have been added to VMM using a

Run As account. VMM automatically uses this Run As account to access the SMB 3.0 file share.

- If you specified explicit user credentials when you added a host or host cluster, you can remove the host or cluster from VMM, and then add it again by using a Run As account.

3. Click **Host Name Properties** > **Storage** > **Add File Share**.
4. In **File share path**, select the required SMB 3.0 file share, and then click **OK**. To confirm that the host has access, open the **Jobs** workspace to view the job status. Or, open the host properties again, and then click the **Storage** tab. Under **File Shares**, click the SMB 3.0 file share. Verify that a green check mark appears next to **Access to file share**.
5. Repeat this procedure for any standalone host that you want to access the SMB 3.0 file share, or for all nodes in a cluster.

Next steps

After you set up the Hyper-V host or cluster, learn about [provisioning VMs](#).

Manage Storage Replica in VMM

12/7/2018 • 6 minutes to read

Storage Replica was introduced in Windows Server 2016. It enables storage-agnostic, block-level, synchronous replication between clusters or servers for disaster preparedness and recovery, as well as stretching of a failover cluster across sites for high availability. Synchronous replication enables mirroring of data in physical sites with crash-consistent volumes, ensuring zero data loss at the file system level. Asynchronous replication allows site extension beyond metropolitan ranges with the possibility of data loss.

[Learn more](#) and review the [FAQ](#).

The article provides explains how Storage Replica integrates with System Center - Virtual Machine Manager (VMM), and describes how to set up Storage Replica using PowerShell to replicate storage in the VMM fabric.

Storage Replica in VMM

You can use Storage Replica to replicate Hyper-V cluster data or file data. Using Storage Replica in VMM provides a number a business advantages:

- Eliminates the cost and complexity associated with synchronous replication solutions such as SAN.
- Synchronous replication minimizes downtime and data loss. It provides a RPO of 0 (zero data loss). RTO (data unavailability) only occurs during the time in which a primary site fails and a secondary site starts.
- Source and destination storage hardware doesn't need to be identical.

Before you start

- VMM must be running on Windows Server 2016 Datacenter Edition.
- Hyper-V must be running on Windows Server 2016 Datacenter, Server Core, or Nano.
- Only synchronous replication is supported. Asynchronous isn't supported.
- You need two sets of storage, either volume or file storage. Both the source and destination locations must have the same type of storage (file or volume) but the actual storage can be mixed. For example you could have Fibre Channel SAN at one end and Spaces Direct (in hyper-converged or disaggregated mode) at the other.
- Each set of storage should be available in each of the clusters. Cluster storage shouldn't be shared.
- Source and destination volumes (including log volumes) need to be identical in size and block size. This is because Storage Replica uses block replication.
- You need at least one 1GbE connection on each storage server, preferably 10GbE, iWARP, or InfiniBand.
- Each file server or cluster node needs firewall rules that allow ICMP, SMB (port 445, plus 5445 for SMB Direct) and WS-MAN (port 5985) bi-directional traffic between all nodes.
- You need to be a member of the Administrator group on each cluster node.
- Storage Replica can only be set up using Windows PowerShell at present.
- Source and destination storage must be managed by the same VMM server.
- Integrating VMM with Azure Site Recovery isn't supported.
- Setting write order and consistency groups isn't supported.

Deployment steps

1. **Identify storage:** Identify the source and destination storage you want to use.
2. **Discover and classify:** If your storage isn't currently in the VMM fabric, you need to discover it with VMM.

Both the source and destination storage must be managed by the same VMM server. After discovery you create a storage pool for it, and a storage classification for it. [Learn more](#).

3. **Pair:** Pair the source and destination storage array.
4. **Provision:** After your storage is paired you'll need to provision identical data and log volumes from the source and destination storage pools created on the respective storage arrays. In addition to provisioning a volume for data that will be replicated, you also need to provision a volume for replication transaction logs. As data is updated on source storage, the transaction log is appended and delta changes are synchronized (using synchronous replication) with destination storage.
5. **Create replication groups:** After the volumes are in place you create replication groups. Replication groups are logical groups containing multiple volumes. The replication groups need to be identical, containing the data and log volumes for the source and destination sites respectively.
6. **Enable replication:** Now you can enable replication between the source and destination replication groups.
7. **Refresh:** To finalize creation of replication groups and to trigger the initial data replication, you need to refresh the primary and secondary storage provider. Data replicates to destination storage.
8. **Verify status:** Now you can check the status of the primary replication group. It should be in the Replicating state.
9. **Add VMs:** When delta replication is up and running you can add VMs that use storage contained in the replication group. When you add the VMs they'll be detected and will begin replicating automatically.
10. **Run failover:** After replication is in a Synchronizing state, you can run a failover to check it's working as expected. There isn't a test failover mechanism right now, so you'll run a manual failover in response to planned or unplanned outages. After failover you can delete the VM on the source site (if it still exists) and create a VM on the destination site using the replicated data.
11. **Run failback:** After failover is complete and replica VMs are up and running, you can fail back as you need to. Note that:
 - If you run an unplanned failover and your source location isn't available, you'll run a failover to failback from the secondary to primary location, and then create the VM in the primary location.
 - If you run a planned failover and the source VM is still available, you need to stop replication, remove the source VM, create the VM in the secondary location, and then restart replication. Then at the primary site you can create the VM with the same settings as the original VM.

Retrieve PowerShell objects

1. Before you start retrieve the name of the PowerShell objects you want to use.
2. Get the name of the primary storage array and assign to variable.

```
$PriArray = Get-SCStorageArray - Name $PriArrayName
```

3. Get the name of the secondary storage array and assign to variable.

```
RecArray = Get-SCStorageArray - Name $RecArrayName
```

4. Get the name of the primary storage pool and assign to variable.

```
$ $ PriPoolName $RecPool = Get-SCStoragePool -Name $
```

5. Get the name of the secondary storage pool and assign to variable.

```
$ $PriPoolName $RecPool = Get-SCStoragePool -Name $
```

Pair the storage arrays

Pair the primary and secondary storage arrays using the variables for the storage array names. Note that the array name should be the same as the cluster name.

```
Set-SCStorageArray -StorageArray $PriArray -PeerStorageArrayName $RecArray.name
```

If you created the cluster outside VMM and you do need to rename the array name to match the cluster name use:

```
Get-SCStorageArray -Name "existing-name" | Set-SCStorageArray -Name "new-name"
```

Provision LUNs and create the storage groups

Provision a LUN from the storage pool for data and for the log. Then create replication groups.

1. Provision and create on the source.

```
Set-SCStorageArray -StorageArray $PriArray -PeerStorageArrayName $RecArray.name

$PrimaryVol = New-SCStorageVolume -StorageArray $PriArray -StoragePool $PriPool -Name PrimaryVol -
SizeInBytes $VolSize -RunAsynchronously -PhysicalDiskRedundancy "1" -FileSystem "CSVFS_NTFS" -DedupMode
"Disabled"

$PrimaryLogVol = New-SCStorageVolume -StorageArray $PriArray -StoragePool $PriPool -Name PrimaryLogVol
-SizeInBytes $LogVolSize -GuidPartitionTable -RunAsynchronously -FileSystem "NTFS"

$PriRG = New-SCReplicationGroup -Name PriRG -StorageVolume $PrimaryVol -LogStorageVolume
$PrimaryLogVol
```

2. Provision and create on the destination.

```
$RecoveryVol = New-SCStorageVolume -StorageArray $RecArray -StoragePool $RecPool -Name RecoveryVol -
SizeInBytes $VolSize -RunAsynchronously -PhysicalDiskRedundancy "1" -FileSystem "CSVFS_NTFS" -DedupMode
"Disabled"

$RecoveryLogVol = New-SCStorageVolume -StorageArray $RecArray -StoragePool $RecPool -Name
RecoveryLogVol -SizeInBytes $LogVolSize -GuidPartitionTable -RunAsynchronously -FileSystem "NTFS"

$RecRG = New-SCReplicationGroup -Name RecRG -CreateOnArray -ProtectionMode Synchronous -StorageVolume
$RecoveryVol -LogStorageVolume $RecoveryLogVol
```

Enable replication

Now enable synchronous replication between the source and destination replication groups.

```
Set-SCReplicationGroup -ReplicationGroup $PriRG -Operation EnableProtection -TargetReplicationGroup $RecRG -
EnableProtectionMode Synchronous
```

Refresh the storage providers

1. Open the VMM console.
2. Click **Fabric Resources** > **Providers**. Right-click the provider > **Refresh**.

Verify replication status

Retrieve the replication status for the source replication group to make sure that replication is working as expected.

```
Get replication status Get-SCReplicationGroup | where {($_.Name.EndsWith("PriRG")) -or
($_.Name.EndsWith("RecRG"))} | fl Name, IsPrimary, ReplicationState, ReplicationHealth
```

Create a VM

Create a VM using a LUN in the source replication group. Alternatively you can create a VM in the VMM console.

```
New-SCVirtualMachine -Name "DemoVM" -VMHost <HostName> -Path $PrimaryVol -VMTemplate <VMTemplate>
```

Run a failover

Run failover.

```
Set-SCReplicationGroup -ReplicationGroup $PriRG -Operation PrepareForFailover

Set-SCReplicationGroup -ReplicationGroup SRecRG -Operation Failover
```

Run failback

Before you fail back, in the VMM console, remove the source VMs if they're still available. You can't fail back to the same VM.

Now run failback:

```
Set-SCReplicationGroup -ReplicationGroup $PriRG -Operation ReverseRoles -EnableProtectionMode Synchronous -
TargetReplicationGroup $RecRG
```

After running failback you can create VMs at the source site using the failed back VHD/configuration files.

Stop replication

If you want to stop replication you'll need to run this cmdlet at the source and destination.

```
Set-SCReplicationGroup -ReplicationGroup $RecRG -Operation TearDown Tear down need to be done on both RGs
```

Learn more

- Learn more about [Storage Replica](#)
- Learn about [allocating storage](#) to Hyper-V hosts and clusters.

Manage scale-out file server (SOFS) in the VMM fabric

12/7/2018 • 2 minutes to read

Scale-out file server (SOFS) is a file server deployed as an active/active cluster based on SMB 3.0. Using an SOFS cluster provides supplies apps with the bandwidth of all nodes in the cluster. All nodes in the cluster accept SMB requests, providing continuous availability and transparent failover if a node goes down.

You can add and manage SOFS clusters in the System Center - Virtual Machine Manager (VMM) fabric. There are a number of ways you can add an SOFS cluster. You can add an existing SOFS cluster to the fabric, provision an SOFS cluster from existing Windows machines in the fabric, or provision a cluster from bare metal computers.

- [Perform a rolling upgrade of an SOFS cluster](#)
- [Add an existing SOFS to the VMM storage fabric](#)
- [Create an SOFS cluster from standalone servers in the VMM fabric](#)
- [Provision SOFS from bare-metal computers](#)

Run a rolling upgrade of an SOFS cluster to Windows Server 2016 in VMM

12/7/2018 • 3 minutes to read

Cluster rolling upgrade was introduced in Windows Server 2016. It enables you to upgrade the operating system of cluster nodes in a scale-out file server (SOFS) cluster, or Hyper-V cluster, without stopping workloads running on the nodes. [Read more](#) about rolling upgrade requirements and architecture.

This article describes how to perform a cluster rolling upgrade of SOFS managed in the System Center - Virtual Machine Manager (VMM) fabric. Here's what the upgrade does:

- **Creates a template:** Creates a template of the node configuration by combining the appropriate physical computer profile with the node configuration settings detailed in the upgrade wizard.
- **Migrates workloads:** Migrates workloads off the node so workload operations aren't interrupted.
- **Removes node:** Puts the node into maintenance mode and then removes it from the cluster. This removes all VMM agents, virtual switch extensions, and so forth from the node.
- **Provisions the node:** Provisions the node running Windows Server 2016, and configures it according to the saved template.
- **Returns the node to VMM:** Brings the node back under VMM management and installs the VMM agent.
- **Returns the node to the cluster:** Adds the node back into the SOFS cluster, brings it out of maintenance mode, and returns virtual machine workloads to it.

Before you start

- The cluster must be managed by VMM.
- The cluster must be running Windows Server 2012 R2.
- The cluster must meet the [requirements](#) for bare metal deployment. The only exception is that the physical computer profile doesn't need to include network or disk configuration details. During the upgrade VMM records the node's network and disk configuration and uses that information instead of the computer profile.
- You can upgrade nodes that weren't originally provisioned using bare metal as long as those nodes meet bare metal requirements such as BMC. You'll need to provide this information in the upgrade wizard.
- The VMM library needs a virtual hard disk configured with Windows Server 2016.

Run the upgrade

1. Click **Fabric > Storage > File Servers**. Right-click the SOFS > **Upgrade Cluster**.
2. In the Upgrade Wizard > **Nodes**, click the nodes you want to upgrade or **Select All**. Then click **Physical computer profile**, and select the profile for the nodes.
3. In **BMC Configuration** select the Run As account with permissions to access the BMC, or create a new account. In **Out-of-band management protocol** click the protocol that the BMCs use. To use DCMI click **IPMI**. DCMI is supported even though it's not listed. Make sure the correct port is listed.
4. In **Deployment Customization**, review the nodes to upgrade. If the wizard couldn't figure out all of the settings it displays a **Missing Settings** alert for the node. For example, if the node wasn't provisioned by bare metal BMC settings might not be complete. Fill in the missing information.
 - Enter the BMC IP address if required. You can also change the node name. Don't clear **Skip Active Directory check for this computer name** unless you're changing the node name, and you want to make sure the new name is not in use.

- In the network adapter configuration you can specify the MAC address. Do this if you're configuring the management adapter for the cluster, and you want to configure it as a virtual network adapter. It's not the MAC address of the BMC. If you choose to specify static IP settings for the adapter select a logical network and an IP subnet if applicable. If the subnet contains an address pool you can select **Obtain an IP address corresponding to the selected subnet**. Otherwise type an IP address within the logical network.
5. In **Summary**, click **Finish** to begin the upgrade. If the wizard finishes, the node upgrades successfully so that all of the SOFS nodes are running Windows Server 2016. The wizard upgrades the cluster functional level to Windows Server 2016.

If you need to update the functional level of a SOFS that was upgraded outside VMM, you can do that by right-clicking the **Files Servers** > SOFS name > **Update Version**. This might be necessary if you upgraded the SOFS nodes before adding it to the VMM fabric, but SOFS is still functioning as a Windows Server 2012 R2 cluster.

Add an existing SOFS to the VMM fabric

12/7/2018 • 2 minutes to read

You can add an existing scale-out file server (SOFS) to the System Center - Virtual Machine Manager (VMM) storage fabric. When you add a file server, VMM automatically discovers all file shares on it.

1. Click **Fabric > Storage > Home > Add Resource > Storage Devices** to open the **Add Storage Devices Wizard**.
2. If the file server is in a domain that's not trusted by the Hyper-V hosts for which it will provide storage, select **This computer is in an untrusted Active Directory domain**. Click **Browse** and select a Run As account with admin permissions on the SOFS. Or create a Run As account if you don't have one.
3. In **Select Provider type**, select **Windows-based file server**.
4. In **Specify Discovery Scope**, type in the FQDN or IP address of the SOFS (not of the underlying cluster)
5. In **Gather Information**, VMM discovers and import information about the SOFS.
6. In **Select Storage Device**, select the file shares that you want VMM to manage. You can configure more settings on the file server after you finish the wizard.
7. On the **Summary** page, confirm the settings and then click **Finish**. You can monitor the cluster status on the **Jobs** page. After the job finishes check the SOFS in **Fabric > Storage > File servers**.

Provision a scale-out file server (SOFS) from standalone file servers in the VMM fabric

12/7/2018 • 2 minutes to read

Use the instructions in this article if you want to use System Center - Virtual Machine Manager (VMM) to create a scale-out-file server (SOFS) from standalone file servers that are managed in the VMM fabric.

1. In the VMM console, click **Fabric > Create > File Server Cluster**.
2. In the **Create Clustered File Server** wizard > **General** specify a cluster name, a file server name, and IP addresses if required.
3. In **Resource Type**, select the option to provision computers on which Windows Server 2012 R2 or later is installed and fill in the details.
4. In **Cluster Nodes**, define a list of computers to add to the cluster.
5. On the **Summary** page, confirm the settings and then click **Finish**.

You can monitor the cluster status on the **Jobs** page. After the job finishes check the cluster in **Fabric > Storage > File servers**.

Provision a scale-out file server (SOFS) cluster from bare metal computers in the VMM fabric

12/7/2018 • 14 minutes to read

In addition to adding existing file servers to an SOFS cluster in the System Center - Virtual Machine Manager (VMM) fabric, VMM can discover provision bare-metal machines as SOFS cluster nodes. This article includes the steps for setting up a bare metal SOFS cluster in VMM.

Before you start

Here's what you need for the deployment:

- **Physical computers** to deploy as SOFS cluster nodes. These computers must meet the prerequisites described in the table below. They can be running on operating system, or an operating system that will be overwritten during the deployment process.
- **Virtual hard disk** with an appropriate operating system, located on a VMM library share. When you create the virtual hard disk, you can create a virtual machine, install the guest operating system, and then use Sysprep with the /generalize and the /oobe options.

The operating system on the virtual hard disk that you deploy on the cluster nodes must support the boot from virtual hard disk (VHD) option.

- **PXE server** configured with Windows Deployment Services is needed for bare metal deployment.

Physical computer requirements

PREREQUISITE	DETAILS
--------------	---------

PREREQUISITE	DETAILS
BMC	<p>Each physical computer must have a baseboard management controller (BMC) installed that enables out-of-band management by VMM.</p> <p>Through a BMC, you can access the computer remotely, independent of the operating system, and control system functions such as the ability to turn the computer off or on.</p> <p>The BMCs must use one of the supported out-of-band management protocols, and the management protocol must be enabled in the BMC settings.</p> <p>Supported protocols: Intelligent Platform Management Interface (IPMI) versions 1.5 or 2.0; Data Center Management Interface (DCMI) version 1.0; System Management Architecture for Server Hardware (SMASH) version 1.0 over WS-Management (WS-Man); custom protocols such as Integrated Lights-Out (iLO)</p> <p>The BMCs should use the latest version of firmware for the BMC model.</p> <p>The BMCs must be configured with logon credentials and must use either static IP addressing or DHCP. If you use DHCP, we recommend that you configure DHCP to assign a constant IP address to each BMC, for example by using DHCP reservations.</p> <p>The VMM management server must be able to access the network segment on which the BMCs are configured.</p>
Operating system	Physical computers must be running Windows Server 2012 R2 or later.
Accounts	<p>You'll need two RunAs accounts.</p> <p>A Run As account for joining computers to the domain, and an account for access to the BMC on each computer.</p>

PFX server requirements

PREREQUISITE	DETAILS
Deployment requirements	<p>You must have a PXE server configured with Windows Deployment Services.</p> <p>If you have an existing PXE server in your environment configured with Windows Deployment Services, you can add that server to VMM. Then you can use it for provisioning in VMM (and VMM will recognize only the resulting servers). All other requests will continue to be handled by the PXE server according to how it is configured.</p> <p>If you don't have an existing PXE server, you can deploy the Windows Deployment Services role on a server running a supported operating system (Windows Server 2012 R2 or later).</p>

PREREQUISITE	DETAILS
Location	The PXE server must be in the same subnet as the physical computers that you want to provision.
Windows Deployment Services installation	<p>When you install Windows Deployment Services you should install both the Deployment server and Transport server options. You don't need to add images.</p> <p>During host deployment, VMM uses a virtual hard disk that you've created and stored in the library.</p> <p>You don't need to configure settings on the PXE response tab. VMM provides its own PXE provider.</p>
Permissions	When you add a PXE server, you must specify account credentials for an account that has local administrator permissions on the PXE server. You can enter a user name and password or specify a Run As account. You can create the RunAs account before you begin, or during deployment.

Virtual disk and template requirements

PREREQUISITE	DETAILS
Virtual hard disk	<p>Make sure you have a generalized virtual hard disk in a VMM library share. It should be running Windows Server 2012 R2 or later.</p> <p>We recommend that for production servers, you use a fixed disk (.vhd or .vhdx file format) to increase performance and to help protect user data.</p> <p>Make sure you have a generalized virtual hard disk in a VMM library share. It should be running Windows Server 2012 R2 or later.</p>
Dynamic disk	When you create a physical computer profile, VMM converts a dynamic disk to a fixed disk.
Custom drivers	If you plan to assign custom drivers to a physical computer profile, you add them to a VMM library share in one or more folders with a .CR (custom resources) extension. VMM recognizes them as custom resources.
Answer file	Like custom resources, if you want a physical computer profile to include references to an answer file (Unattend.xml file), create the answer file and add it to a VMM library share before you start deployment. For example, you might want to create an answer file to enable Remote Desktop Services and place it on a library share. Then you can select that file when you configure a physical computer profile.
RDS	If you use Remote Desktop Services (RDS) to manage servers, we recommend that you enable the RDS connections in the image. You can also enable RDS by using an answer file in the physical computer profile.

PREREQUISITE	DETAILS
Logical networks	<p>If you have already configured logical networks or logical switches in VMM, you can include those configurations in the physical computer profile.</p> <p>To include static IP addressing controlled through a logical network in a physical computer profile configure the logical network. The logical network must include at least one network site and static IP address pool.</p> <p>The network site must also be available to the host group or to a parent host group where you want to assign the hosts that you will be creating from bare metal.</p>
Logical switch	<p>To use a logical switch, install all necessary virtual switch extensions and extension providers, and create the switch before you create the physical computer profile.</p> <p>In the logical switch, as a best practice, include one or more port classifications for the virtual ports.</p> <p>To apply a logical switch to physical adapters in a physical computer profile, make sure that you have installed the intended number of NICs on the physical computer.</p>

Deployment steps

1. **Before you start:** Verify the prerequisites above before you start.
2. **Prepare physical computer:** Set up the BIOS on each physical computer to support virtualization.
3. **Prepare the PXE server environment:** Add the PXE server to the VMM fabric.
4. **Add driver files:** Add driver files to the VMM library if you want to use custom drivers.
5. **Create profile:** Create a profile for the physical computers.
6. **Create the cluster:** Run the Create Clustered File Server Wizard to discover the physical computers, configure the cluster, and start the cluster deployment. The physical computers boot from a customized Windows PE image on the PXE server. The Failover Cluster and File Server roles are enabled. After the cluster is created the Scale-Out File Server role is enabled. The computer is then restarted.
7. **Add nodes to SOFS cluster:** After you've provisioned the nodes you can create a new cluster with them, or add them to an existing one.

Prepare physical computers

Prepare each computer to support virtualization, as follows:

1. Set the BIOS boot order to boot from a Pre-Boot Execution Environment (PXE)-enabled network adapter as the first device.
2. Configure the log on credentials and IP address settings for the BMC on each computer.
3. If your environment has multiple DNS servers, where replication may take some time, we strongly recommend that you create DNS entries for the computer names that will be assigned to the physical computers, and allow time for DNS replication to occur. Otherwise, deployment of the computers may fail.

Add a PXE server to the VMM fabric

1. Click **Fabric > Servers > Home > Add > Add Resources > PXE Server**.

2. In **Computer name** specify the PXE server name.
3. Add the credentials for an account that has local administrator permissions on the PXE server. You can specify an existing Run As account, or create a new account. Manually enter user credentials in the format domain_name\user_name. Then click **Add**.
4. In **Jobs**, verify that the job status is **Completed**, and close the dialog box. The job sets up the new PXE server, installs the VMM agent on the PXE server, imports a new Windows Preinstallation Environment (Windows PE) image, and adds the machine account to VMM for the PXE server.
5. Verify that the PXE server is added in **Fabric > Servers > PXE Servers**. The agent status should be **Responding**.

Add custom resources to the library

If you plan to assign custom drivers, the driver files must exist in the library. You can tag the drivers in the library, so that you can later filter them by tag. After files are added, when you configure a physical computer profile, you can specify the driver files. VMM installs the specified drivers when it installs the operating system on a physical computer.

In the physical computer profile, you can select to filter the drivers by tags, or you can select to filter drivers with matching Plug and Play (PnP) IDs on the physical computer. If you select to filter the drivers by tags, VMM determines the drivers to apply by matching the tags that you assign to the drivers in the library to the tags that you assign in the profile. If you select to filter drivers with matching PnP IDs, you don't need to assign custom tags.

1. Locate a driver package that you want to add to the library.
2. In the library share that is located on the library server that is associated with the group where you want to deploy the physical computers, create a folder to store the drivers, and then copy the driver package to the folder.
3. We strongly recommend that you create a separate folder for each driver package, and that you do not mix resources in the driver folders. If you include other library resources such as .iso images, .vhd files or scripts with an .inf file name extension in the same folder, the VMM library server will not discover those resources. Also, when you delete an .inf driver package from the library, VMM deletes the entire folder where the driver .inf file resides.
4. In the VMM console > **Library > Library Servers**, expand the library server where the share is located, right-click the share, and then click **Refresh**. After the library refreshes, the folder should appear.
5. Assign tags if required. In **Library**, expand the folder that you created to store the drivers, and click the folder that contains the driver package.
6. In the **Physical Library Objects**, right-click the driver .inf file, and then click **Properties**.
7. In the **Driver File Name Properties > Custom tags**, enter custom tags separated by a semi-colon, or click **Select** to assign available tags or to create and assign new ones. If you click **Select**, and then click **New Tag**, you can change the name of the tag after you click **OK**. For example, if you added a network adapter driver file, you could create a tag that is named ServerModel NetworkAdapterModel, where ServerModel is the server model and NetworkAdapterModel is the network adapter model.

Create a physical computer profile

Before you start, determine whether the physical computers use Extensible Firmware Interface (EFI) or BIOS. If you have both, create a separate profile for each type.

1. Click **Library > Home > Create > Physical Computer Profile**.
2. In the **New Physical Computer Profiles Wizard > Profile Description** type in a name and description and select **VM host**.
3. In **OS Image > Virtual hard disk file > Browse**, click the generalized virtual hard disk that you added to the library share. By default, if the disk is dynamic VMM converts it to a fixed disk during host deployment. We

recommend that for production servers you use a fixed disk to increase performance and help protect user data.

4. In **Hardware Configuration > Management NIC** select the network adapter you'll use to communicate with VMM and whether to use DHCP or a static address. If you want to use Consistent Device Naming (CDN) for the adapter or configure logical switches and ports click **Physical Properties**. Click **Add** to add the adapter.
5. In **Disk** specify the partitioning scheme for the first disk. You can use GPT if the physical computer profile is EFI. In **Partition Information** select the volume label, whether to use all remaining free space or a specific size, and whether to designate the partition as the boot partition. You can also add a new disk or partition. During deployment VMM will copy the virtual hard disk file to the boot partition and automatically create a system partition on the same disk.
6. In **Driver filter** filter the drivers that will be applied to the operating system during host deployment. You can filter by Plug and Play ID or by specific tags. If you select to filter drivers with matching tags make sure you've added driver files to the library and assigned corresponding tags.
7. In **OS Configuration** specify the domain that the Hyper-V host or cluster should join, specify local admin credentials and identity information. Add the product key for installation, and set the time zone. In **GUIRunOnce** you can specify one or more commands that will run when the user logs on to the Hyper-V host for the first time.
8. In **Host Settings** specify the path of the host to store the files that are associated with virtual machines placed on the host. Don't specify drive C because it's not available for placement. If you don't specify a path VMM placement will determine the most suitable location.
9. In **Summary** verify the settings. Wait until Jobs shows a status of completed, and then verify the profile in **Library > Profiles > Physical Computer Profiles**.

Provision a Scale-Out File Server cluster from bare metal

The Create Clustered File Server Wizard does the following:

1. Discovers the physical computers through out-of-band management
2. Deploys the Windows Server operating system image on the computers by using the physical computer profile (if configured to do so)
3. Enables the file server role on the computers
4. Enables the Scale-Out File Server role on the cluster
5. Adds the provisioned computers as a Scale-Out File Server cluster under VMM management

Run the wizard:

1. Click **Fabric > Servers > Home > Create > File Server Cluster**.
2. In the **Create Clustered File Server Wizard > General** type in a cluster name, file server name, and cluster IP addresses if needed.
3. In **Resource Type** select the option to provision bare-metal computers. Select the physical computer profile and click **Next**.
4. In **Credentials and Protocols** click **Browse** next to the Run As account and choose the account with permissions to access the BMC. In the **Protocol** list click the out-of-band management protocol you want to use for discovery. If you want to use DCMI click **Intelligent Platform Management Interface (IPMI)**. DCMI 1.0 isn't listed but it is supported. Make sure you use the latest version of firmware for the BMC model.
5. In **Discovery Scope** specify the IP address scope that includes the IP addresses of the BMCs. You can add a single address, a subnet, or range.
6. In **Target Resources** select the computers you want to provision, allow time for deep discovery, and click

items to review and modify information. Note that if the number of physical network adapters doesn't match the number of physical adapters defined in the computer profile you'll need to add the missing information. If you don't want to deploy a computer immediately you can select its BMC IP address and click **Remove**.

7. In **Deployment Customization** configure the settings and when there are no more warnings about missing information click **Next**.
 - **DHCP**: If your physical computer profile uses DHCP click an BMC IP address and type in a computer name. Decide whether to skip the AD check. If you do the check deployment will continue if the computer account exists. Click the entry for each BMC IP address.
 - **Static**: If the profile uses static IP addresses for each BMC IP address, type in a MAC address of the computer's network adapter that's used to communicate with VMM. Click the logical network you want to use. The default logical network is the one indicated in the profile. Click the IP subnet you want to use. The subnet list is scope to what's defined for the logical network in the associated network sites. You should select the IP subnet that corresponds to the physical location in which you're deploying the server and the network to which the adapter is connected. you can automatically assign an IP address or assign a specific address.
8. In **Summary**, confirm the settings and click **Finish**. To confirm the cluster was added, click **Fabric > Storage > File Servers**.

Manage SOFS settings in the VMM fabric

5/6/2019 • 4 minutes to read

You can manage scale-out file server (SOFS) in the System Center - Virtual Machine Manager (VMM) fabric as follows:

- **Create storage pools:** Create storage pools from physical disks on SOFS nodes. You can then allocate those storage pools.
- **Create file shares:** You can create file shares on a SOFS in the VMM fabric. You can set the storage type for a share as a storage pool, local pool, or volume.
- **QoS:** Set up a quality-of-service (QoS) policy for SOFS to control resources allocated to VMs
- **Set a disk witness for a storage pool:** From VMM 2016, you can specify that the disk witness for an SOFS cluster should come from a particular storage pool. To do this VMM creates a three-way mirror space and configures it as the disk witness for the cluster. The storage pool must have at least five physical disks.

Create storage pools

1. Click **Fabric > Storage > File Servers**. Right-click the SOFS server (not the nodes) and click **Manage Pools**.
2. In **Manage Pools of File Server**, click **New** to create a new pool or modify an existing one.
3. In **General**, specify a name and select a storage classification.
4. In **Physical Disks**, select the disks you want to include in the pool. Disks will be displayed in accordance with the SOFS settings. For example, a file server with shared storage might show SAS storage disks, or a file server using Storage Spaces Direct would show local disks attached to each node.
5. In **Default Settings**, leave the default unless you need to change for a specific reason.
 - **Fault domain:** When you select a fault domain you specify how many copies of your data will be distributed across a cluster. Note that fault domain isn't displayed for clusters configured with Storage Spaces Direct. These clusters have a fault domain of **Node** to indicate that copies of data are stored on multiple nodes in the cluster and data is available even if a specific node isn't.
 - **Interleave:** Interleave (along with the number of columns) specify the way in which data is written to physical disks.
6. Click **OK** to save the storage pool settings. After the job completes verify the pools in **Fabric > Storage > Classifications and Pools**.

Create a file share

1. Click **Fabric > Storage > Home > Create File Share**.
2. In **Create File Share Wizard > Storage Type**, select the SOFS on which you want to create a share. Type in a name and description for the share and select the storage pool you want to use. If the CSV exists, click **Volume** and specify it. If the folder path exists click **Local path** and specify it. The file share inherits the classification of the storage pool.
3. In **Capacity**, specify the file share size and type. Leave the default type unless the disk is used for backups or deduplication, in which case NTFS is recommended.
4. In **Capacity > Resiliency**, for ReFS resilience must be mirror (two or three-way). For NTFS, it can be mirror or parity (single or dual). The default is a three-way mirror.

5. Enable deduplication if necessary. Change the unit size allocation if required, and optionally enable storage tiers.
6. In **Summary**, review the settings and click **Finish**. Verify the file share in **Fabric > Storage > File Servers > File Shares**.

Set a storage QoS for an SOFS

System Center VMM 2016 and later include storage QoS policies to solve the "noisy neighbor" problem. This problem is a common in virtualized environments. When two virtual machines (VMs) share a resource, say a disk, there is always a chance that one VM's usage of the resource exceeds that of the other. This can impact app performance of an app running on the VM. Storage QoS provides:

- **Mitigation of noisy neighbor issues:** Ensures that a single VM doesn't consume all resources and starve other VMs of storage bandwidth.
- **Monitor end-to-end storage performance:** When VMs are started on an SOFS their performance is monitored.
- **Manage storage I/O in accordance with business needs:** Storage QoS policies define minimums and maximums for VMs and ensure they are met, even in over-provisioned environments. If policies can't be met alerts are issued.

Set a storage QoS policy as follows:

1. Click **Fabric > Storage > Storage. QoS Policies > Create Storage Qos Policy**.
2. In **Create Storage Qos Policy Wizard > General**, specify a name and description for the policy.
3. In **Policy Settings** select whether you want all the virtual hard disks for VMs to share resources equally or allocate resources per VM. If you choose to allocate per instance you'll need to set a minimum and maximum IOPS. Then a virtual disk to which the policy is applied will receive the minimum and maximum limits.
4. In **Scope** specify the file servers on which to apply the policy. You can apply the policy to multiple servers which is useful when you migrate VMs across servers so that QoS policy settings remain the same.
5. In **Summary** review the settings and click **Finish**. Verify the policy in **Fabric > Storage > QoS Policies**.

Set a disk witness for the SOFS

1. Click **Fabric > Storage > File Servers**. Right-click the SOFS server (not the nodes) and click **Properties**.
2. In **General** click **Use disk witness for this file server from the specified pool**, to indicate that the disk witness for the SOFS should come from a specific storage pool. VMM creates a three-way mirror space, and configures it as the disk witness for the cluster.

Provision virtual machines in the VMM fabric

3/14/2019 • 13 minutes to read

This article provides an overview of provisioning Virtual Machines (VMs) in the System Center - Virtual Machine Manager (VMM) compute fabric. Learn about provisioning methods, and features provided by VMM during provisioning.

Provisioning

VMs can be provisioned using a number of methods:

- [Create VMs from a blank virtual hard disk](#): You create a VM and install an operating system from an .iso image, removable media, or from a network boot with a PXE server.
- [Create a VM from an existing virtual hard disk](#): You create a VM from a virtual hard disk in the VMM library. We recommend a VHD that's been generalized with Sysprep.
- [Clone a VM from an existing VM](#): Clone an existing VM in the VMM library to create a new one. We recommend you clone a VM that's been generalized with Sysprep.
- [Create a VM from a template](#): Create VMs with consistent settings configured in a VM template. VM templates are XML objects stored in the VMM library. They can be used to control and restrict VM settings available to self-service users. Template settings include the guest operating system profile, a hardware profile, and one or more VHDs that can be used to create a new VM.
- **Create a VM in a service deployment**: In VMM you can create services that are logical grouping of VMs that are configured and deployed as a single entity. A single tier service includes a single VM. Multi-tier services have multiple VMs.
- [Rapidly provision a VM using storage area network \(SAN\) copy](#): Deploy a VM using SAN copy abilities such as snapshot and clone. You can rapidly provision standalone VMs, or VMs that are provisioned in a service.

Deploying a VM guest cluster

On a Hyper-V cluster, you can deploy a guest failover cluster that consists of multiple VMs that uses shared .vhdx files. VMM supports the following:

You can deploy a guest failover cluster that uses shared .vhdx files on a Hyper-V failover cluster. In this scenario, if Hyper-V uses Cluster Shared Volumes (CSV) on block-level storage, then the shared vhdx files are stored on a CSV that's configured as shared storage. Alternatively, Hyper-V can use SMB file-based storage deployed by Scale-Out File Server (SOFS), as the location of the shared .vhdx files. No other storage types are supported, and third-party SMB storage isn't supported.

VM placement

When you deploy or migrate a VM, VMM uses intelligent VM placement to evaluate available hosts.

- The placement algorithm analyzes performance data for the workload and the host, and then rates hosts on a scale of one to five stars to indicate the best placement choice.
- Placement includes a preferred and possible owners feature that allows to specify which hosts are preferred and possible if failover of VMs occurs.
- Placement considers storage classifications. Clouds can be scope to limit VM placement to specific storage classifications only

- Placement options can be selected as follows:
 - **Create a new VM:** The placement process offers a suggestion for the host. If a self-service user creates a VM, the host is automatically assigned by VMM, based on the highest rating. **Migrate a VM:** During migration VMM provides host ratings to help you select a host. **Convert a VM to Hyper-V:** The conversion wizard provides rating for hosts so that you can select the best one.

Host ratings

- VMM evaluates all hosts within a selected host group and any hosts contained in child host groups. Host ratings are calculated on a scale from 0 to 5 stars, where 5 stars indicate the highest rating. The ratings are based on default criteria that don't include all information. For example, network connection speed is not taken into account.
- Ratings are based on individual hosts, and not on the relative suitability of all available hosts. Ratings for one host do not change based on the ratings for other hosts.
- VMM calculates host ratings according to specific formulas, described in the following table.
 - **CPU rating:** $[1 - (\text{CPU Utilization} / (100 - \text{CPU Reserve}))] \times \text{CPU Weight}$
 - **Memory (RAM) rating:** $[1 - (\text{Memory Used} / (\text{Total Memory} - \text{Memory Reserve}))] \times \text{Memory Weight}$
 - **Disk I/O capacity rating:** $[1 - (\text{Disk IOPS} / \text{Maximum Disk IOPS})] \times \text{Disk Weight}$
 - **Network rating:** $[1 - (\text{Network Utilization} / (100 - \text{Network Reserve}))] \times \text{Network Weight}$
- A host is rated only when a virtual machine needs to be placed.
- The information gathered about a host is compared to the information about the resources required by the VM, and a rating is assigned to the host.
- During automatic placement, VMM attempts to use the host assigned the highest rating. During manual placement the host rating is shown so that you can select the appropriate host. As of System Center 2012 SP1, you can select a host in VMM even if not all hosts have been rated. The selected host must have a positive number of stars.
- VMM measures CPU, memory, disk, and network usage approximately every 10 minutes to recalculate an average rating that is an average of all the measurements taken that the last action that reset the host rating. Host ratings are reset when the following happens:
 - A new virtual machine is created
 - A virtual machine is deployed, stored, migrated or deleted
 - A virtual machine is turned on, off, or moved into a stopped, paused, or saved state.

Zero rating

A host might be assigned a zero rating if it does not meet conditions to receive a non-zero rating. To receive a non-zero rating the following criterion are required:

- The host must have at least one hard disk with enough storage to hold the total hard disk space required by the virtual machine. With dynamic hard disks, the current hard disk size is used, not the maximum hard disk size.
- The memory required for the virtual machine must be less than the current amount of memory available on the host. A host must also have sufficient memory available to start the virtual machine. Note that VMM does offer the option of overcommitting cloud and host group capacity for replica VMs.
- If dynamic memory is enabled note the following:
 - If the virtual machine (including any one of its checkpoints) is configured to use Dynamic Memory, the host should also have Dynamic Memory enabled. If it does not the placement of the virtual machine

- will be blocked during creation or migration.
- For placement of a new or stopped virtual machine, the host must meet at least the startup memory requirement for the virtual machine.
- For placement of a running virtual machine the host must meet at least the current memory requirement for the virtual machine.
- For placement of a virtual machine in a saved state, the last known memory usage value of the virtual machine will be compared to the startup memory of the virtual machine.
- The host must contain all of the virtual networks required for the virtual machine. If you use network tags, the network location tags for the virtual machine and host must be identical.
- A host in maintenance mode automatically receives a zero rating.
- If Microsoft RemoteFX 3D video adapter is enabled on the virtual machine, the host must support RemoteFX and have one or more RemoteFX-capable graphics processing units (GPUs) with sufficient available memory. If these conditions aren't available and the virtual machine is running, placement will be blocked. If it is stopped or in a saved state, a zero rating with a warning will be issued, but placement will not be blocked.
- Highly available virtual machines must be placed on clustered hosts. VMM assigns zero stars to hosts that are not clustered but manual placement is not blocked. If you migrate a highly-available virtual machine to a non-clustered host, the virtual machine will no longer be highly available after the migration.
- VMM blocks migration of Hyper-V hosts to hosts running different virtualization software. Migration of a virtual machine with specific features not allowed by the virtualization software that is running on a host will be blocked. For example, Hyper-V hosts do not allow booting up from a SCSI hard disk.

Handling generation 1 and 2 VMs

In VMM you can create generation 1 and generation 2 VMs.

- You can add VM templates specifying either generation 1 or generation 2 to a service template. Properties will appear for the generation you choose.
- Generation 2 VMs can only run on a host with a host operating system that supports them - Windows Server 2012 R2 onwards. For example, VMM won't place a generation 2 VM on a host running Windows Server 2012.
- When you use a virtual hard disk in .vhd format for a VM or VM template, the VM is automatically generation 1 because .vhd doesn't support generation 2. If you use .vhdx you can select which option you want to use.
- If you use an existing virtual machine or virtual machine template as the starting point for a new virtual machine or virtual machine template, the generation is determined by the existing virtual machine or virtual machine template.
- If you create a hardware profile you can choose between generation 1 and 2. When you incorporate the profile into a VM or VM template you specify the generation you want to use.
- In all wizards and PowerShell cmdlets, you'll be offered configuration options in line with the generation of the VM.
- Generation 1 and 2 VMs don't handle boot order in the same way.
 - You can customize the start order for generation 1 VMs in the VMM console, in the hardware settings when you create a VM. You can also customize using the BootOrder PowerShell parameter.
 - To customize the start order for generation 2 VMs you need to use PowerShell, using the FirstBootDevice parameter when you create a VM. For example, to set the first boot device as the network adapter type: **Set-SCVMTemplate -Template "Generation2template" -FirstBootDevice "NIC,0"**

Optimizing provisioning

Differencing disks

- A differencing disk is a VHD that stores changes made about another VHD or guest operating system by storing them in a separate file.
- When you create a differencing disk you associate another VHD with it (the parent disk). The differencing disk stores all the changes that would otherwise be made to the parent disk if the differencing disk didn't exist.
- In essence the differencing disk saves changes without altering the parent disk.
- Multiple differencing disks can use the same parent VHD.

VMM optimizes support for differencing disks to provide the following:

- Optimized migration of storage that utilizes differencing disks. During a migration, VMM does not migrate base disks unless it is necessary.
- Optimize virtual machine deployment time by utilizing differencing disks. VMM will attempt to identify and utilize differencing disks on the target computer.
- When differencing disks is utilized, deployment of the base virtual disk is optimized by taking advantage of the Windows Offloaded Data Transfers (ODX) capability to copy files to the guest machine during service deployment.
- Optimize time and storage of cloning of virtual machines by utilizing differencing disks. VMM provides the option to create and utilize differencing disks during a cloning operation.
- Note that if parent disks are lost or corrupted all VHDs that depend on it are lost too. You should have a backup plan in place.

If you're using differencing disks you should optimize management by making sure that unused parent HDs are removed regularly.

Fast file copy

During virtual machine deployment, VMM needs to move and copy large files such as VHDs between two locations.

Fast file copy improves the performance of file transfers, mostly by leveraging [Windows Offloaded Data Transfers \(ODX\)](#). In VMM, background intelligent transfer (BITS) is still used as a mechanism for file transfers, but VMM leverages ODX when possible (for example when copying files to SANs that support ODX). This greatly improves the time performance of virtual machine deployments.

Provisioning VMware VMs

After you've set up a vCenter server and ESX/ESXi hosts, you can provision VMware VMs like any other VM. Note the following:

- You can organize and store VMware VMs in the VMM library.
- You can organize and store VMware virtual machines, .vmdk (VMDK) files, and VMware templates in the VMM library.
- You can create VMs from VMware templates stored in the library. You can also import templates stored on vSphere hosts (only template metadata is imported to VMM). VMM doesn't support older VMDK disk types. These disk types are supported:
 - Regular VMDK files (VMFS and monolithic flat)
 - VMDK files that are used to access physical disks (vmfsPassthroughRawDeviceMap)
 - Snapshots (vmfssparse).
- You can create new VMware VMs from VMDK templates.
- You can deploy VMM services to vSphere hosts, but not vApps.

- You can place VMs on VMware hosts based on host ratings when you create, deploy, or migrate VMware VMs. This includes concurrent VM deployment when you're deploying a service.
- You can use dynamic optimization and power optimization for VMware VMs.
- You can do regular VMM networking tasks, including assigning logical networks, static IP address and MAC address to Windows-based VMs running on VMware.
- VMM doesn't support VMware VMs with VHDs connected to an IDE bus.
- VMM supports VMware thin provision hard disk through the dynamic disk type. Note that if you create and deploy a VM to a vSphere host configured to use a dynamic disk the disk will be thin provisioned. If a VM was created as a thin provisioned disk out-of-band VM will display it as dynamic. If you save a thin provision disk to the library VMM will save it as thick. It remains thick if you create a VM from it.

Converting VMware VMs to Hyper-V

There are currently a couple of methods for converting VMWare VMs to Hyper-V (V2V):

- **Convert Virtual Machine Wizard:** In the VMM console you can use this wizard. This method has a number of limitations:
 - Supported for vSphere 4.1 onwards.
 - You can't convert VMware workstations.
 - You can't convert VMs with virtual hard disks connected to an IDE bus.
 - Online conversions aren't supported. You need to power off the VMware VMs.
 - Anti-virus apps must be supported.
 - VMware tools must be uninstalled from the guest operating system of the VM.
- [Microsoft Virtual Machine Converter](#): After June 3, 2017 this tool will no longer be available. [Learn more](#)

Static MAC address for VMs deployed on a VMM cloud

In earlier releases, VMM allowed the users to set a static MAC address on the VMs deployed on the hosts, and did not have an option to set static MAC address for the VMs deployed on the cloud.

VMM 2019 allows you to set a static MAC address for the VMs deployed on VMM Cloud.

Use the following steps:

NOTE

MAC address that you wish to assign to the VM should be part of an accessible MAC pool. As self-service users do not have visibility into the fabric MAC pools, they would need to co-ordinate with admins to make sure that the MAC address is part of the accessible MAC pool.

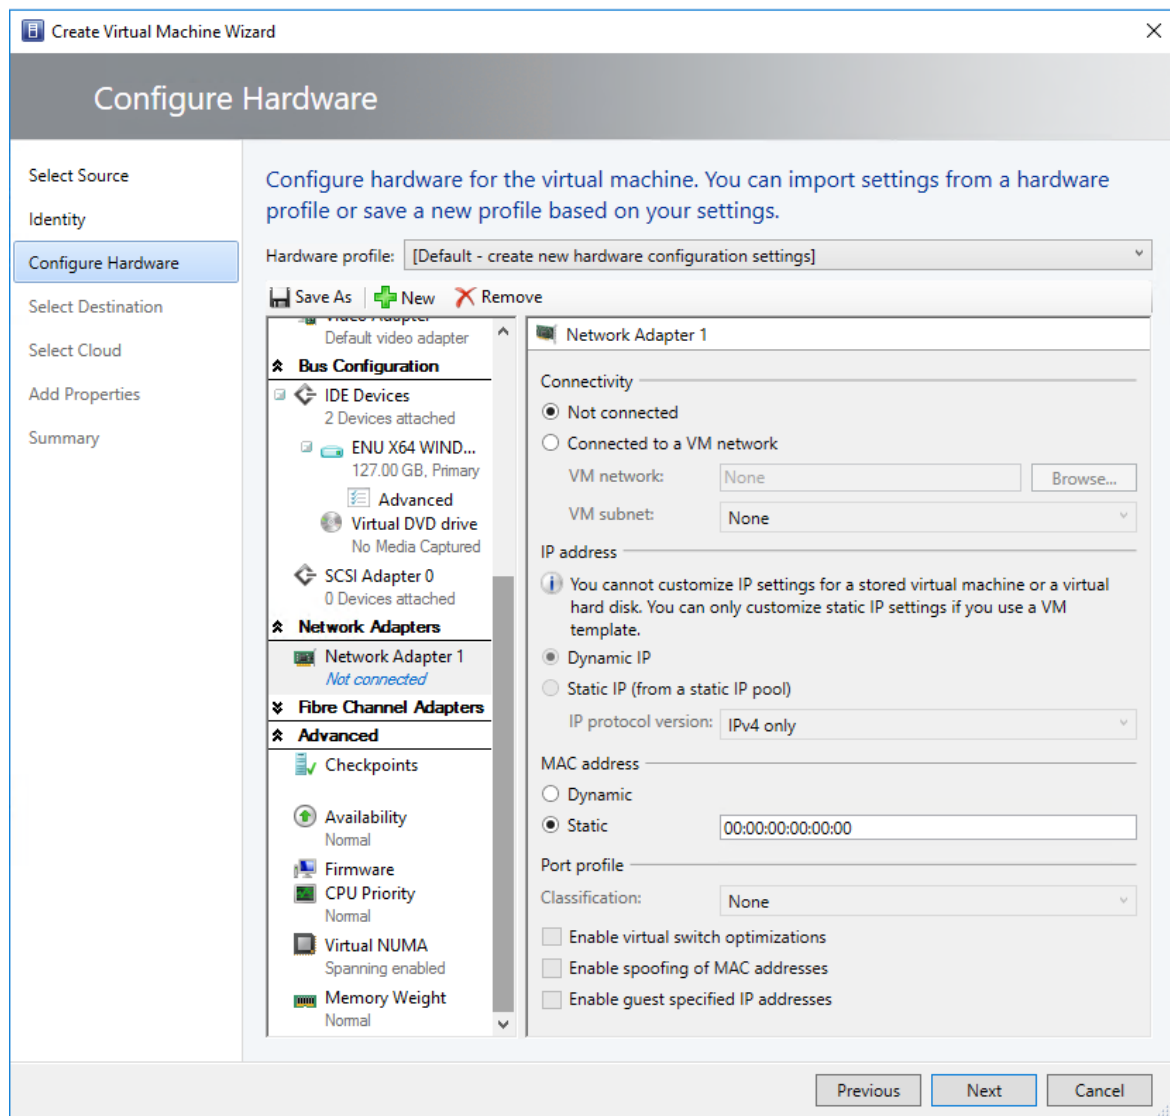
You can set the static MAC address on the VM while:

- Deploying a new VM onto the cloud from VHD/VM Template. or
- Changing the MAC address on an existing VM deployed to the cloud.

Set static MAC while deploying a new VM onto the cloud from VHD/VM template.

1. In VMM Console, navigate to **VMs and Services > Home > Clouds**.
2. Select the cloud to deploy the VM.
3. Right-click > **Create Virtual Machine** to enter the Create Virtual Machine wizard and create the VM.
4. Right-click the VM, select **Properties**.

5. Navigate to **Configure Hardware** page, select the network adapter to which you want to assign a static MAC address.
6. In the **MAC address** section, select **Static** and enter the MAC address in the text box.



Change the MAC address for a VM deployed on the cloud

1. Navigate to a VM deployed on the cloud, for which you want to assign a static MAC address.
2. Follow the steps 4 to 6 from the above procedure.

Next steps

- [Deploy VMs from a blank virtual hard disk](#)

Deploy VMs in the VMM fabric from a blank virtual hard disk

12/7/2018 • 6 minutes to read

This article describes how to create and deploy virtual machines in the System Center - Virtual Machine Manager (VMM) fabric from a virtual hard disk (VHD).

Before you start

- To complete the steps you must be an Administrator or Delegated Administrator on the VMM server, or a self-service user.
- If you're a self-service user you need **Deploy** permissions with the **Store and re-deploy** action assigned. You must first deploy the VM to a private cloud, and then store it in the library.
- You can only customize static IP address settings if you create a VM from a VM template.
- You can use VMM to configure the availability settings for the virtual machine. [Learn more](#)

Create a VM

1. Click **VMs and Services** > **Create Virtual Machine** > **Create Virtual Machine**.
2. In **Create Virtual Machine Wizard** > **Select Source**, click **Create the new virtual machine with a blank virtual hard disk** > **Next**.
3. In **Identity**, specify the VM name and an optional description. In the **Generation** box, select **Generation 1** or **Generation 2**. Then click **Next**.
4. In **Configure Hardware** page, either select the profile that you want to use from the **Hardware profile** list, or configure the hardware settings manually. The hardware setting displayed will differ depending on whether you're deploying a generation 1 or generation 2 machine. Then click **Next**.
 - In **Compatibility**, if you want to deploy the virtual machine to a private cloud, select a capability profile that is available to the private cloud.
 - In **Bus Configuration**, if you want to install an operating system from a DVD or an .iso image, ensure there is a virtual DVD drive that is configured to use an available option such as the **Existing ISO image file** option. If you want to use an ISO image file, the file must be present in the VMM library.
 - If you want to store the virtual machine in the VMM library before you deploy it to a host, use one of the blank virtual hard disks that are provided by default in the VMM library. Click the VHD in **Bus Configuration**. Click **Use an existing virtual hard disk** > **Browse** and select a blank hard disk
 - If the virtual machine is a generation 1 that boots from the network to install an operating system, in **Network Adapters**, use the legacy network adapter type.
5. In **Select Destination** page, specify how the virtual machine should be deployed - in a private cloud, on a host, or stored in the library.

Deploy the VM in a private cloud

1. In **Select Cloud** select the private cloud on which you want to place the virtual machine. If you are

connected as an Administrator, you can select the host on which the virtual machine should be deployed in the private cloud. Cloud suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:

- **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation** or **Deployment and Transfer Explanation:** Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).
 - The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
 3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
 4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
 5. In **Summary**, confirm the settings, and click **Create**. Confirm that the VM was created in **VMs and Services > Clouds**, and select the cloud. The virtual machine appears in the **VMs** pane.

Deploy the VM on a host

1. In **Select Host**, view the ratings, click the host on which you want to deploy the VM, and click **Next**. The host suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation** or **Deployment and Transfer Explanation:** Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).

- The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more.](#)
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
 3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
 4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
 5. On the **Summary** page, confirm the settings, and then click **Create**.

Store the VM in the library

1. In **Select Library Server**, click the library server that you want to use, and then click **Next**.
2. In **Select Path**, specify the library share location to store the virtual machine. Click **Browse** to select a library share and an optional folder location, click **OK**, and then click **Next**.
3. In **Summary** confirm the settings, and then click **Create**.
4. To confirm that the virtual machine was created, in the **Library** workspace, in the **Library** pane, expand **Library Servers**, expand the library server where you stored the virtual machine, and then click **Stored Virtual Machines and Services**. The stored virtual machine appears in the **Physical Library Objects** pane.

Next steps

After you create the virtual machine, you can install an operating system from an .iso image, from a CD or DVD, or from a network boot if a Pre-Boot Execution Environment (PXE) server is available.

Deploy VMs in the VMM fabric from an existing virtual hard disk

12/20/2018 • 6 minutes to read

This article describes how to create and deploy virtual machines in the System Center - Virtual Machine Manager (VMM) fabric, from an existing virtual hard disk (VHD).

Before you start

- To complete the steps you must be an Administrator or Delegated Administrator on the VMM server, or a self-service user.
- If you're a self-service user you need **Deploy** permissions with the **Store and re-deploy** action assigned. You must first deploy the VM to a private cloud, and then store it in the library.
- The VHD you want to use must be stored in the VMM library. [Learn more](#)
- Use a VHD that has been generalized using Sysprep. If you don't the identity of the cloned VM will be the same as the source VM, and issues might occur if you turn them on together.

Create a VM

1. Click **VMs and Services > Create Virtual Machine > Create Virtual Machine**.
2. In **Create Virtual Machine Wizard > Select Source**, click **Use an existing virtual machine, VM template, or virtual hard disk > Browse**. Select an existing VHD.
3. In **Identity**, specify the VM name and an optional description. If the VHD you choose is in the .vhdx format, in the **Generation** box, select **Generation 1** or **Generation 2**. Then click **Next**.
4. In **Configure Hardware**, either select the profile that you want to use from the **Hardware profile** list, or configure the hardware settings manually. The hardware setting displayed will differ depending on whether you're deploying a generation 1 or generation 2 machine. Then click **Next**.
 - In **Compatibility**, if you want to deploy the virtual machine to a private cloud, select a capability profile that is available to the private cloud.
 - In **Bus Configuration**, if you want to install an operating system from a DVD or an .iso image, ensure there is a virtual DVD drive that is configured to use an available option such as the **Existing ISO image file** option. If you want to use an ISO image file, the file must be present in the VMM library.
 - If you want to store the virtual machine in the VMM library before you deploy it to a host, use one of the blank virtual hard disks that are provided by default in the VMM library. Click the VHD in **Bus Configuration**. Click **Use an existing virtual hard disk > Browse** and select a blank hard disk
 - If the virtual machine is a generation 1 that boots from the network to install an operating system, in **Network Adapters**, use the legacy network adapter type.
5. In **Select Destination** page, specify how the virtual machine should be deployed - in a private cloud, on a host, or stored in the library.

Deploy the VM in a private cloud

1. In **Select Cloud** select the private cloud on which you want to place the virtual machine. If you are connected as an Administrator, you can select the host on which the virtual machine should be deployed in the private cloud. Cloud suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation** or **Deployment and Transfer Explanation:** Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).
 - The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
5. In **Summary**, confirm the settings, and click **Create**. Confirm that the VM was created in **VMs and Services > Clouds**, and select the cloud. The virtual machine appears in the **VMs** pane.

Deploy the VM on a host

1. In **Select Host**, view the ratings, click the host on which you want to deploy the VM, and click **Next**. The host suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation** or **Deployment and Transfer Explanation:** Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the

same VMM library, library share, and logical unit number (LUN).

- The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
 3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
 4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
 5. In **Summary**, confirm the settings, and then click **Create**.

Store the VM in the library

1. In **Select Library Server**, click the library server that you want to use, and then click **Next**.
2. In **Select Path**, specify the library share location to store the virtual machine. Click **Browse** to select a library share and an optional folder location, click **OK**, and then click **Next**.
3. In **Summary** confirm the settings, and then click **Create**.
4. To confirm that the virtual machine was created, in the **Library** workspace, in the **Library** pane, expand **Library Servers**, expand the library server where you stored the virtual machine, and then click **Stored Virtual Machines and Services**. The stored virtual machine appears in the **Physical Library Objects** pane.

Next steps

- [Manage the VM settings](#)

Clone VMs from another VM in the VMM fabric

12/7/2018 • 6 minutes to read

This article describes how to create and deploy virtual machines by cloning existing VMs in the System Center - Virtual Machine Manager (VMM) fabric.

Before you start

- To complete the steps you must be an Administrator or Delegated Administrator on the VMM server, or a self-service user.
- When you clone a virtual machine, the existing virtual machine source is not deleted. We recommend that you clone a virtual machine that has been prepared and generalized with the Sysprep tool. If you do not use a generalized virtual hard disk, the identity of the new virtual machine will be the same as the source. Issues might occur if you turn on two virtual machines with the same identity at the same time.
- You can clone a virtual machine that is deployed on a host. The machine can be online, stopped, or in a saved state.
- You can clone a VM that is stored in the VMM library.
- The option to use differencing disk optimizations is automatically applied when you deploy the virtual machine on a host, if a base disk exists on that host.
- If you're a self-service user you need **Deploy** permissions with the **Store and re-deploy** action assigned. You must first deploy the VM to a private cloud, and then store it in the library.

Create a VM

1. Click **VMs and Services > Create Virtual Machine > Create Virtual Machine**.
2. In **Create Virtual Machine Wizard > Select Source**, click **Use an existing virtual machine, VM template, or virtual hard disk > Browse**. Select an existing VM.
3. In **Select Source**, click **Next**.
4. In **Configure Hardware**, optionally configure any available settings. Then click **Next**.
5. In **Select Destination** page, specify how the virtual machine should be deployed - in a private cloud, on a host, or stored in the library.

Deploy the VM in a private cloud

1. In **Select Cloud** select the private cloud on which you want to place the virtual machine. If you are connected as an Administrator, you can select the host on which the virtual machine should be deployed in the private cloud. Cloud suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation or Deployment and Transfer Explanation:** Lists any factors that make a storage

area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).

- The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
 3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
 4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
 5. In **Summary**, confirm the settings, and click **Create**. Confirm that the VM was created in **VMs and Services > Clouds**, and select the cloud. The virtual machine appears in the **VMs** pane.

Deploy the VM on a host

1. In **Select Host**, view the ratings, click the host on which you want to deploy the VM, and click **Next**. The host suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization**: Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available**: With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details**: Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation**: Provides an explanation if a host received a zero rating.
 - **SAN Explanation** or **Deployment and Transfer Explanation**: Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).
 - The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
5. On the **Summary** page, confirm the settings, and then click **Create**.

Store the VM in the library

1. In **Select Library Server**, click the library server that you want to use, and then click **Next**.
2. In **Select Path**, specify the library share location to store the virtual machine. Click **Browse** to select a library share and an optional folder location, click **OK**, and then click **Next**.
3. In **Summary** confirm the settings, and then click **Create**.
4. To confirm that the virtual machine was created, in the **Library** workspace, in the **Library** pane, expand **Library Servers**, expand the library server where you stored the virtual machine, and then click **Stored Virtual Machines and Services**. The stored virtual machine appears in the **Physical Library Objects** pane.

Next steps

- [Manage the VM settings](#)

Deploy VMs with rapid provisioning using SAN copy in the VMM fabric

12/7/2018 • 11 minutes to read

This article describes how to rapidly provision VMs in the System Center - Virtual Machine Manager (VMM) fabric using SAN copy.

Rapid provisioning provides a method for deploying new virtual machines to storage arrays without needing to copy VMs over the network. VMM uses the SAN infrastructure for cloning VMs, with a VM template to customize the guest operating system.

- You can use rapid provisioning to deploy standalone VMs, and VMs that are deployed as part of a service.
- You create a SAN copy-capable template from a virtual hard disk (VHD) that resides on a storage logical unit that supports SAN copy through cloning or snapshots.
- When you create a VM by using the SAN copy-capable template, VMM quickly creates a read-write copy of the logical unit that contains the VHD, and places the virtual machine files on the new logical unit. When VMM deploys a virtual machine by using rapid provisioning through SAN copy, VMM uses a SAN transfer instead of a network transfer. During a SAN transfer, a SAN copy of the logical unit that contains the virtual machine is created and is assigned to the destination host or host cluster. Because the files for a virtual machine are not actually moved over the network when you transfer a virtual machine over a SAN, it is much faster than a transfer over a standard network.
- You can use either of the following methods to create a SAN copy-capable template.
 - Create a SAN-copy capable template from a new VM
 - Create a SAN-copy capable template from an existing VM

Before you start

- Any storage that is accessible by the provisioned computer may be partitioned during the provisioning process even if a specific disk is selected to be used as the operating system disk. In this case data will be lost. To guarantee the use of a specific boot volume, use deep discovery and do not restart the computer before the deployment of the operating system completes.
- The storage array must support the new storage management features in VMM.
- The storage array must support cloning or snapshots, and the cloning or snapshots feature must be enabled. This might require additional licensing from your storage vendor.
- The storage pool that you want to use for rapid provisioning must be under VMM management. To meet this requirement, you must add the Storage Management Initiative Specification (SMI-S) provider for the array, discover storage pools, classify the storage, and set the preferred allocation method for the storage array to either snapshot or cloning.
- The storage pool that you want to use for rapid provisioning must be allocated to the host group where you want to use rapid provisioning of virtual machines.
- The Hyper-V hosts that you want to use as placement destinations must be members of the host group. Additionally, the following prerequisites must be met:
 - If you want to create a SAN-copy capable template from a new virtual machine, the host where you create the virtual machine must also be a member of this host group.
 - If you want to create a SAN-copy capable template from an existing virtual machine, and want to create and assign the logical unit from the library server, the library server must be a member of this host group. Therefore, the library server must be a Hyper-V host. (If you do not want to add the library server

as a host, you can assign the logical unit out-of-band by using your storage array vendor's management tools.)

- If you want to use rapid provisioning to deploy generation 2 VMs, you must choose a host with an operating system that supports them.
- All Hyper-V hosts that you want to use for rapid provisioning and the library server must have access to the storage array. Also, they must use the same type of SAN connectivity. For SAN migrations to succeed, you cannot have some hosts that connect to the array through Fibre Channel and others that connect through iSCSI. Configuration varies, depending on your storage hardware.
- You should get specific configuration information from the storage vendor, but configuration typically requires:
 - The Multipath I/O (MPIO) feature must be added on each host that will access the Fibre Channel or iSCSI storage array. You can add the MPIO feature through Server Manager.
 - If the MPIO feature is already enabled before you add a host to VMM management, VMM will automatically enable MPIO for supported storage arrays by using the Microsoft provided Device Specific Module (DSM). If you already installed vendor-specific DSMs for supported storage arrays, and then add the host to VMM management, the vendor-specific MPIO settings will be used to communicate with those arrays.
 - If you add a host to VMM before you add the MPIO feature, you must manually configure MPIO to add the discovered device hardware IDs. Alternatively you can install vendor-specific DSMs.
 - If you are using a Fibre Channel storage area network (SAN), each host that will access the storage array must have a host bus adapter (HBA) installed. Additionally, ensure that the hosts are zoned accordingly so that they can access the storage array.
 - If you use an iSCSI SAN, ensure that iSCSI portals have been added and that the iSCSI initiator is logged into the array. Additionally, ensure that the Microsoft iSCSI Initiator Service on each host is started and set to Automatic. For information about how to create an iSCSI session on a host through VMM, see [How to Configure Storage on a Hyper-V Host in VMM](#).

Create a SAN copy-capable template from a new virtual machine

Create a new VM on a logical unit assigned to a Hyper-V host. On the library server, create a SAN-copy capable template from the VM. Note that:

- The library server doesn't need to be a managed Hyper-V host, but it must be able to access the storage pool in which the logical unit resides.
 - When you create the template the logical unit is automatically unregistered from the host, and registered to the library server.
1. Create a logical unit in the VMM storage fabric, from the managed storage pool you want to use for rapid provisioning. Alternatively you can create an assign the logical unit in your storage array management tool.
 2. Allocate the logical unit to the host group where the target host resides. Then assign to logical unit to the host. When you assign the LUN you can format it and assign a drive letter. Note that the logical unit you want to assign must be empty.
 3. Create a virtual machine with a blank virtual hard disk file on the logical unit.
 - In **Select Source**, select **Create the new virtual machine with a blank virtual hard disk**.
 - In **Configure Hardware**, configure the required settings. Make sure **Create a new virtual hard disk** is selected.
 - In **Select Destination**, accept the default setting to **Place the virtual machine on a host**.
 - In **Configure Settings**, in **Select Destination Folder**, click the drive that you created from the assigned logical unit. Verify that **SAN (Migration Capable)** appears next to the drive information. For example: **(L:) [9.92 GB free of 10.00 GB, SAN (Migration Capable)]**.
 - In **Machine Resources**, click **Virtual Hard Disk**. In **Browse > Select Destination Folder**, click the

drive you created from the assigned logical unit.

- In **Select Network** and **Add Properties**, select the required settings. In **Summary**, review the settings and click **Create**. Verify that the VM is listed in **VMs and Services > All Hosts > VMs**.
4. On the new VM, install and customize the guest operating system and the applications that you want. Generalize the image by using Sysprep.exe with the **/generalize** and the **/oobe** options to generalize the associated virtual hard disk. [Learn more](#). When you're finished, make sure there are no .iso image files attached to the virtual DVD drive.

Create a SAN copy-capable template from an existing VM

Create a template from an existing VM.

- If you want to perform this procedure in VMM, the library server must be added as a managed Hyper-V host. This enables you to assign the logical unit to the library server through VMM. If you do not want to make the library a managed Hyper-V host, you can use your array vendor's management tools to assign the logical unit to the library server.
 - You must have an existing virtual hard disk (that was generalized by using Sysprep) that you want to use as a base image for rapid provisioning.
 - Create a folder in the library share that you will use to mount the logical unit to, and to store the virtual hard disk. For example, create a folder in the SEALibrary library share that is named Rapid Provision VHD.
1. Create a logical unit in the VMM storage fabric, from the managed storage pool you want to use for rapid provisioning.
 2. Format the logical unit, and mount it to the folder path you created.
 3. Assign the logical unit to the library server. If the library server is a managed Hyper-V host, you can create and assign the logical unit from the library server. You can also format the disk with NTFS and mount the logical unit to the folder path in the library share at the same time.
 - When you create the logical unit, select the option **Mount in the following empty NTFS folder > Browse**, and then click the folder that you created.
 - Do not assign a drive letter. Also, do not ever create multiple mount points to the folder.
 4. If the library server isn't a managed Hyper-V host, use your array vendor's management tools to create the logical unit, and to unmask the logical unit to the library server. Then do the following:
 - Don't assign a drive letter.
 - Use Disk Management (diskmgmt.msc) to rescan the disk, initialize the disk, and then format it.
 - in Disk Management, mount the logical unit to the folder path you created in the library share (**Change Drive Letter and Paths > Add > Mount in the following empty NTFS folder**, and click the empty library folder).
 5. Copy the virtual hard disk you want to use to the new folder in the library share. Note that the virtual hard disk should be the only file on the logical unit.
 6. The new folder that you created appears in the library share. To verify the virtual hard disk SAN copy-capable, click the new folder, and in **Physical Library Objects**, click the VHD file. **SAN copy capable** should indicate **Yes**.

Create a SAN-copy capable template

1. Click **Library > Create > Create VM Template**.
2. In **Create VM Template Wizard > Select Source**, click **From an existing virtual machine that is deployed on a host > Browse**. Select the VM on the logical unit. Click **Yes** on the warning message.
3. In **Identity** type in a template name and description.

4. In **Configure Hardware** click **Next**. The classification that appears matches what you assigned to the storage pool from which you created the logical unit.
5. In **Configure Operating System**, click **Next**.
6. In **Select Library Server**, click the library server where you want to create the template. Verify that the **Transfer Type** is **SAN**, and then click **Next**. The library server must have access to the same storage pool as the host.
7. In **Select Path**, click **Browse**, and select a location on the library server to store the VM files.
8. In **Summary** review the settings and click **Create**. In **Jobs** you can track the template being created. Wait for the **Completed** status. Verify the template in **Library > Templates > VM Templates**.

Deploy a VM from the template

New deploy a VM from the SAN-copy capable template. This procedure explains how to deploy a standalone VM. Alternatively you can select the template when you [create a service](#). Note that:

- The hosts where you want to place the VMs must have access to the managed storage pool where the logical unit that is associated with the template resides.
 - If you want to deploy the virtual machines to a private cloud, the storage classification that is assigned to the logical unit that was used to create the SAN clone-capable template must be available to the private cloud.
 - For cloud deployment, the host groups that are used to provide resources for the private cloud must contain the hosts that have access to the managed storage pool where the logical unit that is associated with the template resides.
1. Click **VMs and Services > Create > Create Virtual Machine**.
 2. In the Create Virtual Machine Wizard > **Select Source**, click **Use an existing virtual machine, VM template or virtual hard disk > Browse**. Select type **VM Template**, and click the template you created for rapid provisioning. The template should indicate **Yes** in the **SAN Copy Capable** column.
 3. In **Select Source**, click **Next**.
 4. Complete the rest of the steps wizard to create and deploy the virtual machine. Note the following:
 - In **Configure Hardware > Bus Configuration**, leave the **Classification** list empty, or select the storage classification that - In **Select Host** or **Select Cloud**, make sure that the **Transfer Type** column indicates **SAN**. If you selected to place the virtual machine on a host, in **Configure Settings > Machine Resources**, click the virtual hard disk to verify the deployment options. For rapid provisioning through SAN copy, make sure that the method to deploy the virtual hard disk to the host list is **Transfer the virtual disk by using the SAN**.
 5. After you complete the wizard, open **Jobs > Create virtual machine job** to view the job status.
 6. When you create a virtual machine from the SAN copy-capable template, a new logical unit is automatically provisioned from the same storage pool where the virtual hard disk that was used to create the SAN copy-capable template from resides. The logical unit is automatically registered and mounted on the target host.
 7. To verify that the virtual machine was created, open the VMs and Services workspace. Expand **All Hosts or Clouds**, and locate and click the destination host or private cloud. In **VMs**, verify that the new virtual machine appears. If you open Disk Management (Diskmgmt.msc) on the destination host, you can see the new disk that is assigned and registered to the host.

Next steps

[Manage the VM settings](#)

Create and deploy VMs in the VMM fabric from a VM template

12/7/2018 • 7 minutes to read

This article describes how to create VMs in the System Center - Virtual Machine Manager (VMM) fabric from a VMM virtual machine (VM) template. You can use a VM template to create standalone VMs, or to create VMs in tiers in a [service template](#).

Before you start

- Some settings, including server roles and features, application installation, and SQL Server settings, apply only when a VM template is used for service deployments. For standalone virtual machine creation, these settings are not used and won't appear when you create a standalone VM.
- The ability to configure a VM to use static IP addresses from an IP address pool managed by VMM is only available when you deploy a VM using a VM template.
- To complete the steps you must be an Administrator or Delegated Administrator on the VMM server, or a self-service user.
- If you're a self-service user you need **Deploy** permissions with the **Store and re-deploy** action assigned. You must first deploy the VM to a private cloud, and then store it in the library.
- You can only customize static IP address settings if you create a VM from a VM template.
- You can use VMM to configure the availability settings for the virtual machine. [Learn more](#).

Create a VM

1. Click **VMs and Services** > **Create Virtual Machine** > **Create Virtual Machine**.
2. In **Create Virtual Machine Wizard** > **Select Source**, click **Use an existing virtual machine, VM template, or virtual hard disk** > **Browse**.
3. In **Select Virtual Machine Source**, select the template > **OK**. Click **OK** if a message appears that some deployment settings will be ignored.
4. In **Select Source**, click **Next**.
5. In **Identity**, specify the VM name and an optional description. Then click **Next**.
6. In **Configure Hardware** page, either select the profile that you want to use from the **Hardware profile** list, or configure the hardware settings manually. The hardware setting displayed will differ depending on whether you're deploying a generation 1 or generation 2 machine. Then click **Next**.
 - In **Compatibility**, if you want to deploy the virtual machine to a private cloud, select a capability profile that is available to the private cloud.
 - In **Network Adapters**:
 - If you want to use static IP addresses, set the MAC address to static.
 - If the VM uses a VHD in VMware .vmdk format, include a legacy network adapter in the template (**New** > **Network Adapter** > **Legacy Network Adapter**. If you don't the VM might not be able to start in a domain, although it's OK in a workgroup.
 - In **Configure Operating System**, specify the guest operating system settings. If you have an existing profile select in the **Guest OS profile** list.

7. In **Select Destination** page, specify how the virtual machine should be deployed - in a private cloud, on a host, or stored in the library.

Deploy the VM in a private cloud

1. In **Select Cloud** select the private cloud on which you want to place the virtual machine. If you are connected as an Administrator, you can select the host on which the virtual machine should be deployed in the private cloud. Cloud suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.
 - **SAN Explanation or Deployment and Transfer Explanation:** Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).
 - The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
5. In **Summary**, confirm the settings, and click **Create**. Confirm that the VM was created in **VMs and Services > Clouds**, and select the cloud. The virtual machine appears in the **VMs** pane.

Deploy the VM on a host

1. In **Select Host**, view the ratings, click the host on which you want to deploy the VM, and click **Next**. The host suggestions are based on a 0-5 star rating. [Learn more](#). Verify the settings and modify if required:
 - **Expected utilization:** Expected utilization for a VM created from a blank VHD is based on standard defaults. VMM updates host suggestions and ratings in response to modifications made to the expected virtual machine utilization.
 - **Make this VM highly available:** With this option selected, only hosts that are located in a cluster are available for selection.
 - **Details:** Indicates the status of the host, the operating system, and the type and status of the virtualization software.
 - **Rating Explanation:** Provides an explanation if a host received a zero rating.

- **SAN Explanation** or **Deployment and Transfer Explanation**: Lists any factors that make a storage area network (SAN) transfer unavailable. VMM does not recognize a virtual machine that is stored on a SAN as available for deployment by using SAN transfer if the virtual machine was stored directly in the library when it was created or was added to the library during a library refresh. To avoid this issue, deploy the virtual machine to a host by using a LAN transfer, and then store the virtual machine in the same VMM library, library share, and logical unit number (LUN).
 - The **Deployment and Transfer Explanation** tab provides an explanation if fast file copy cannot be used. Fast file copy is based on the Windows Offloaded Data Transfers (ODX). [Learn more](#).
2. In **Configure Settings**, review the VM settings. Either accept the default VM path on the host or specify a different location. you can optionally select to **Add this path to the list of default virtual machine paths on the host**. In **Machine Resources** accept the default values for the VHD, or modify as required. To prevent placement from choosing its own values, click the pin icon next to the setting. This option isn't available for self-service users.
 3. In **Select Networks**, if it appears, optionally select the network settings, and click **Next**.
 4. In **Add Properties**, configure the action to take when the host starts or stops, and the operating system that you will install on the VM. Then click **Next**.
 5. On the **Summary** page, confirm the settings, and then click **Create**.

Store the VM in the library

1. In **Select Library Server**, click the library server that you want to use, and then click **Next**.
2. In **Select Path**, specify the library share location to store the virtual machine. Click **Browse** to select a library share and an optional folder location, click **OK**, and then click **Next**.
3. In **Summary** confirm the settings, and then click **Create**.
4. To confirm that the virtual machine was created, in the **Library** workspace, in the **Library** pane, expand **Library Servers**, expand the library server where you stored the virtual machine, and then click **Stored Virtual Machines and Services**. The stored virtual machine appears in the **Physical Library Objects** pane.

Next steps

- [Configure the VM settings](#)

Create and deploy Linux virtual machines in the VMM fabric

3/14/2019 • 2 minutes to read

This article describes how to create and deploy Linux VMs in the System Center - Virtual Machine Manager (VMM) fabric.

Before you start

VMM supports virtual machines that contain Linux as the guest operating system. Note that:

- Linux Integration Services (LIS) must be installed on the virtual machine.
- The VMM guest agent for Linux must be installed on the virtual machine. It is required for service template integration, and it allows you to modify properties on the Linux computer such as the host name.

NOTE

SCVMMLinuxGuestAgent (XPlat) cfghostdomain garbles the hosts file if hostname is already set to localhost. We recommend you not to set the hostname as localhost when deploying Linux VMs through VMM.

- VMM doesn't verify that the VM meets these requirements. However, if it doesn't, VM deployment will fail.

Create the VM

Create a VM running Linux using any of the available methods in the VMM fabric. [Learn more](#).

Install LIS on the VM

By default, LIS is included with some distributions of Linux. If LIS is not included in the distribution of Linux that you are using for the virtual machine, then install it manually. [Learn more](#).

Install the VMM guest agent

1. Open an elevated command prompt on the VMM server.
2. Go to the **c:\Program Files\Microsoft System Center 2012\Virtual Machine Manager\agents\Linux** folder.
3. Copy all the agent installation files from that folder to a new folder on the VM.
4. Open the new folder on the VM, and run the following command: **chmod +x install**.
5. Run either of these commands, depending on the operating system.

```
./install scvmmguestagent.1.0.0.544.x64.tar  
./install scvmmguestagent.1.0.0.544.x86.tar
```

When the agent installs on the VM the following files and folders will be created on the VHD:

- A default installation folder (/opt/microsoft/scvmmguestagent), and an installation log file (scvmm-install.log)

- A default log files folder - /var/opt/microsoft/scvmmagent/log
- A specialization log file (scvmm.log). This file is created when the virtual machine is deployed and specialized.
- A configuration file (scvmm.conf). This file contains the location of the log file and is used to control logging during deployment and specialization.

Next steps

[vm-settings](#)

Configure a nested VM as a host

7/9/2019 • 2 minutes to read

Nested virtualization is a functionality in Windows Server 2016 and above that allows you to run Hyper-V inside a Hyper-V virtual machine. In other words, with nested virtualization, a Hyper-V host itself can be virtualized. Nested virtualization can be enabled out-of-band by using PowerShell and Hyper-V host configuration.

You can leverage this functionality to reduce your infrastructure expense for development and test scenarios without the need for individual hardware.

System Center - Virtual Machine Manager (VMM) 1801 and later, you can enable and disable the nested virtualization feature through VMM console. You can configure the nested Virtual Machine (VM) as a host in VMM and perform host operations from VMM, on this VM. For example, VMM dynamic optimization will consider a nested VM host for placement.

NOTE

Virtualization applications other than Hyper-V are not supported in Hyper-V virtual machines, and are likely to fail. This includes any software that requires hardware virtualization extensions.

Before you start

Ensure the following prerequisites are met:

- A Hyper-V host running Windows Server 2016, Windows Server 2019.
- A Hyper-V VM running Windows Server 2016, Windows Server 2019.
- A Hyper-V VM with configuration version 8.0 or greater.
- An Intel processor with VT-x and EPT technology.

Ensure the following prerequisites are met:

- A Hyper-V host running Windows Server 2016.
- A Hyper-V VM running Windows Server 2016.
- A Hyper-V VM with configuration version 8.0 or greater.
- An Intel processor with VT-x and EPT technology.


Procedure - configure a nested VM as a host

Enable the nested virtualization on a VM and then configure it as a host. Administrators/delegated administrators can configure nested virtualization by using VMM. Use the following two procedures:

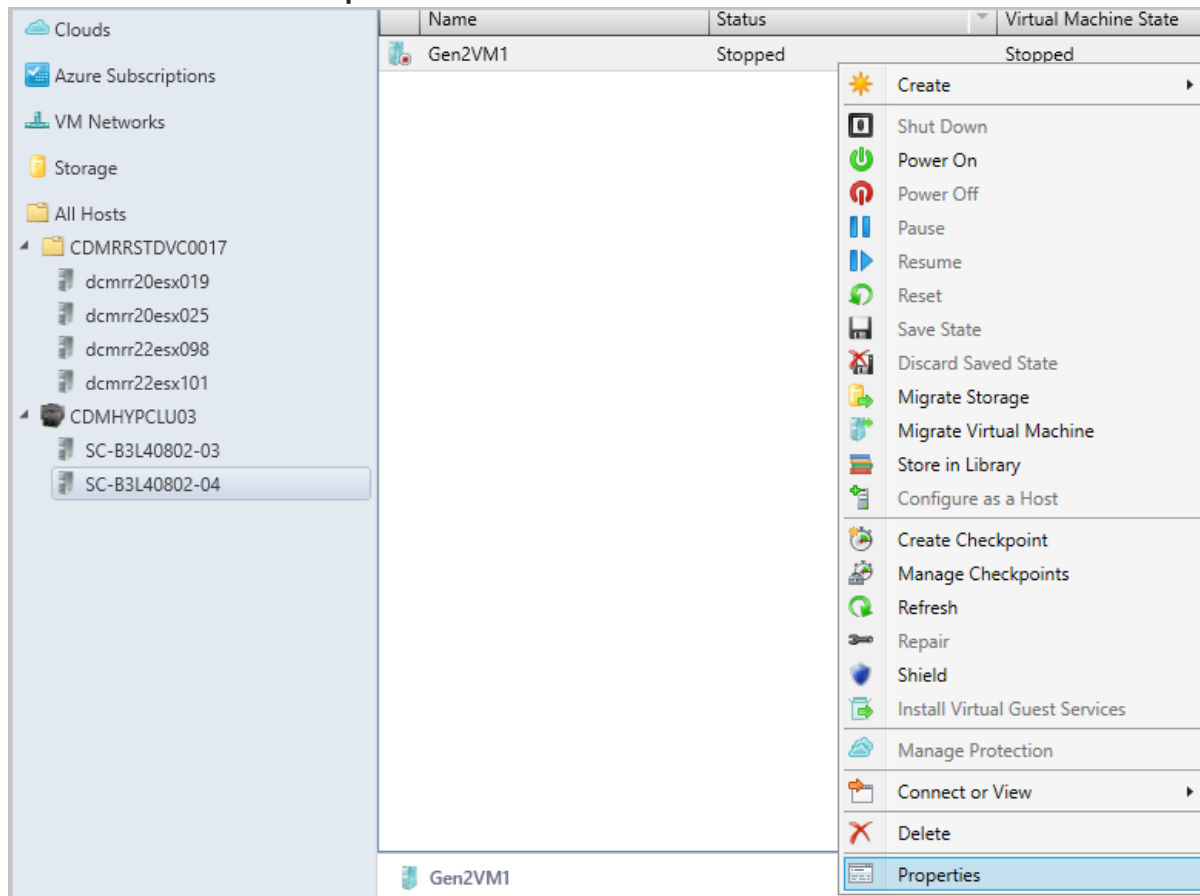
- [Enable nested virtualization on a VM](#)
- [Configure the VM as a host in VMM](#)

Enable nested virtualization on an existing virtual machine

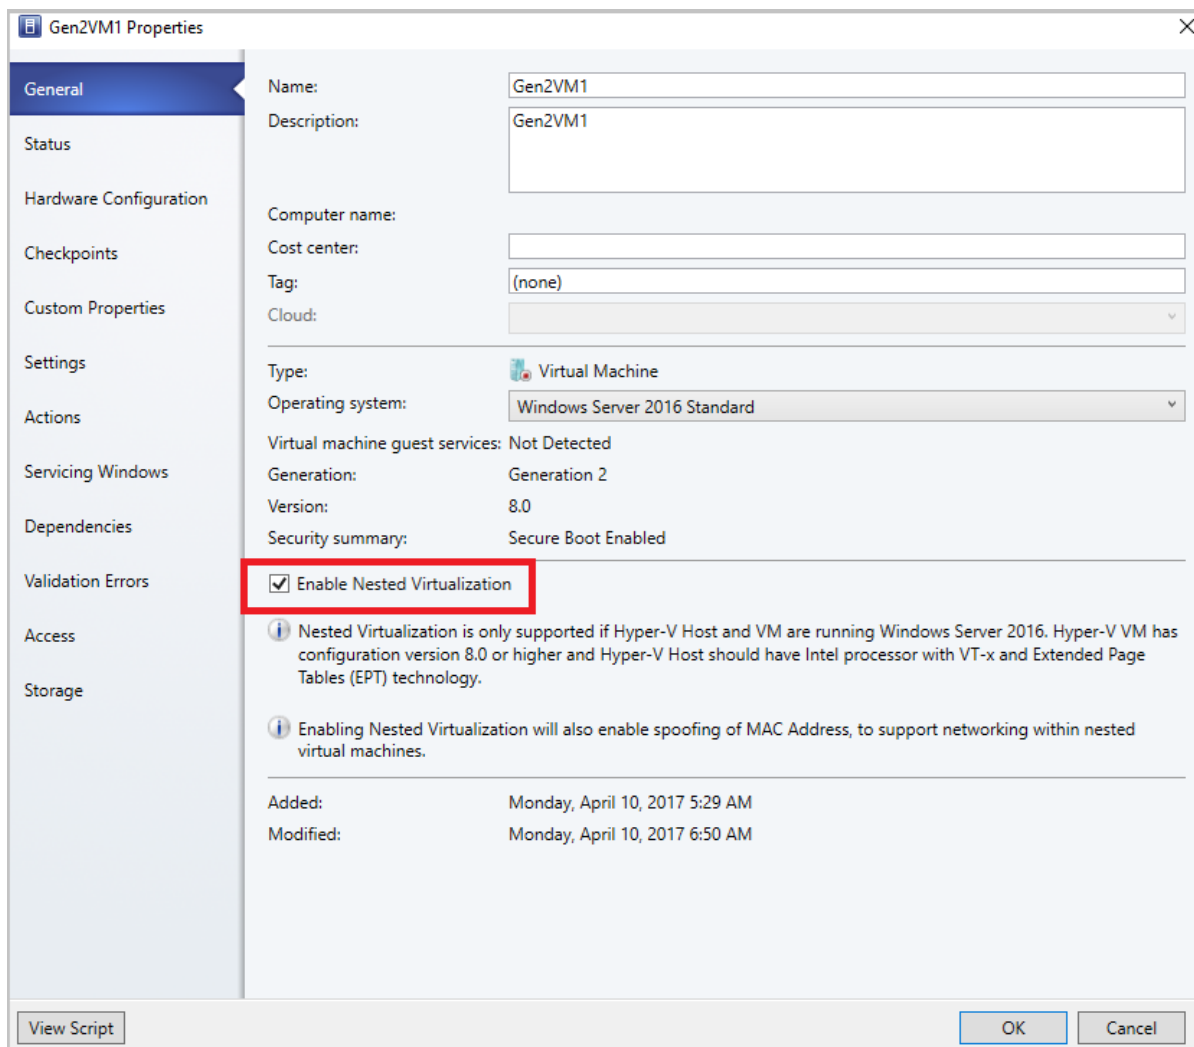
1. Identify the VM that meets the above [prerequisites](#).
2. Ensure the VM is in **stopped** state.



3. Browse the selected VM's **Properties**.



4. On **General**, select **Enable Nested Virtualization**.



Configure the nested VM as a host

1. Enable the following inbound and outbound firewall rules on the nested VM that you want to configured as the host.

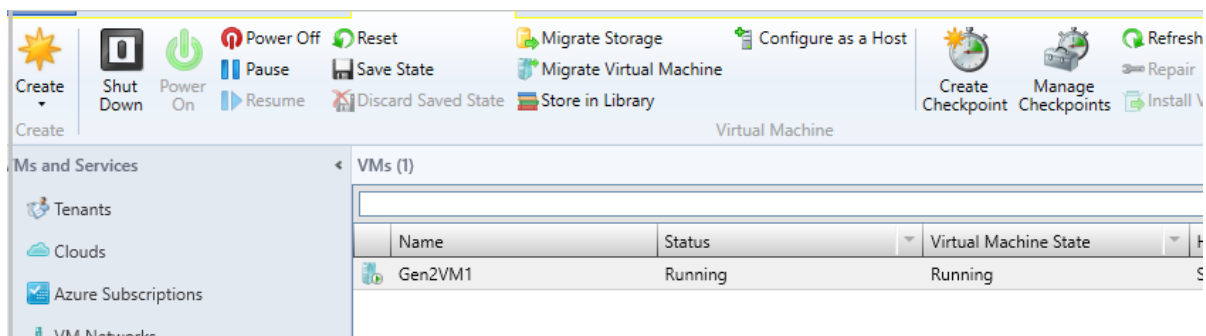
Inbound Firewall rules

- File and printer sharing
- Windows remote management (HTTP-In)
- Windows management Instrumentation

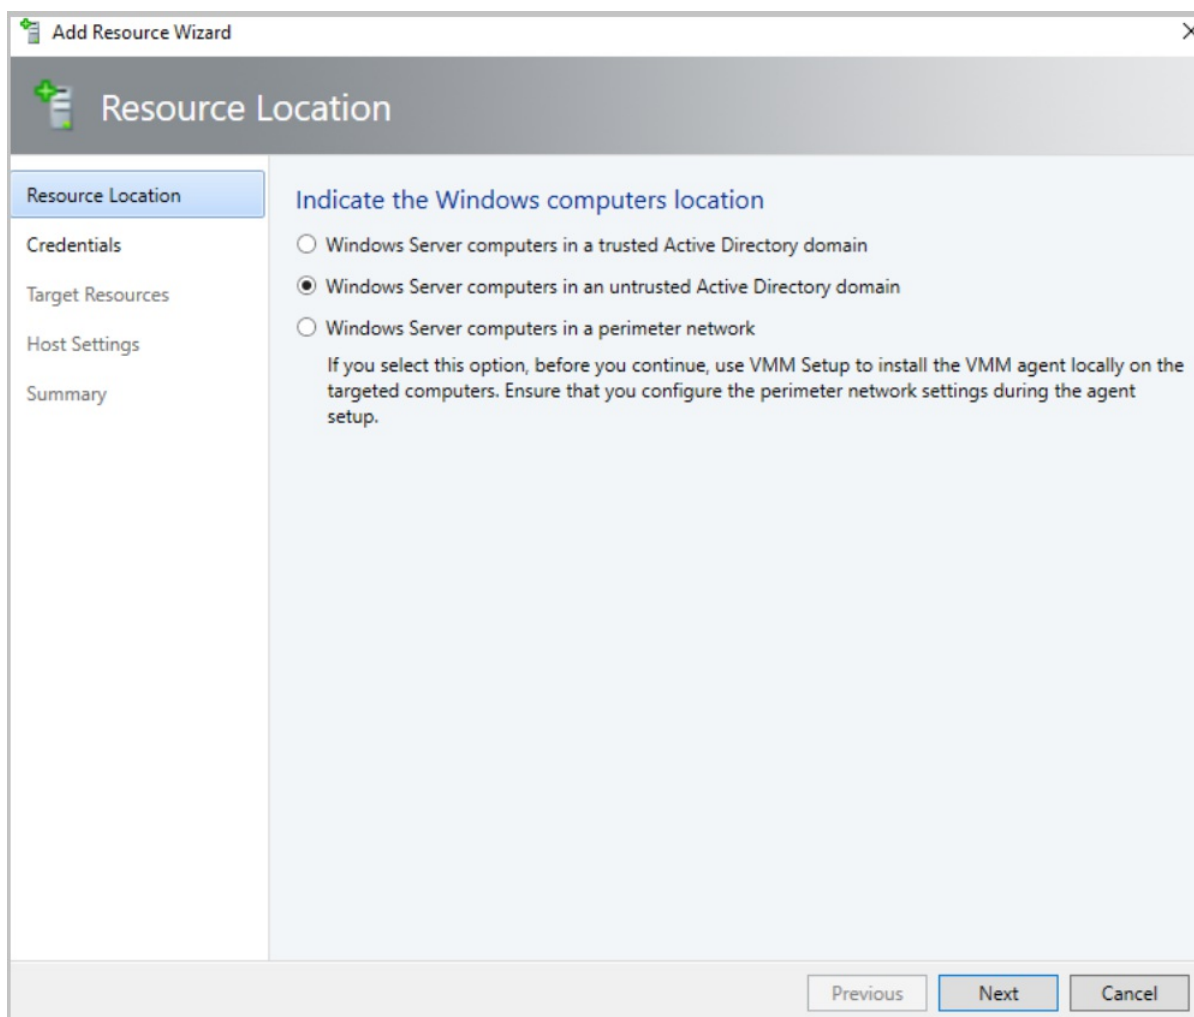
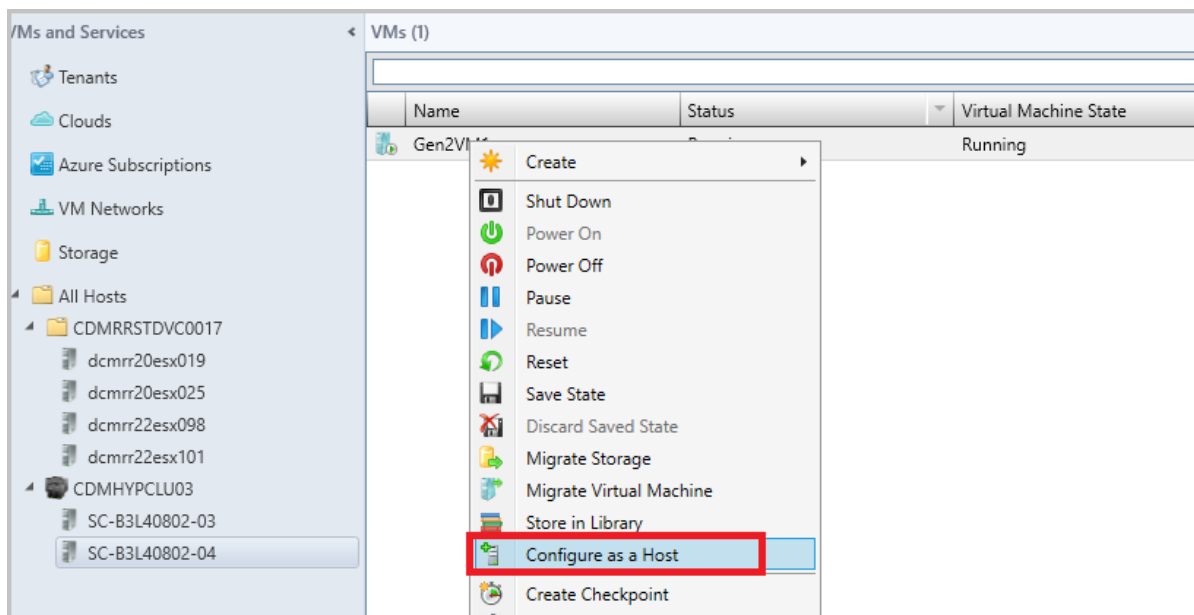
Outbound Firewall rules

- File and printer sharing
- Windows management instrumentation (WMI-Out)

2. Ensure the VM is in **running** state. Start the VM if it is not running.



3. Right-click the VM and select **Configure as a Host**. The **Add Resource** wizard appears.



4. Run through the wizard, select the options as appropriate and complete the wizard.

Disable nested virtualization

1. Select the host or VM for which nested virtualization is enabled.
2. Ensure the VM is in **stopped** state. Stop the VM if it is running.
3. Browse the VM **Properties**.
4. On **General**, clear the **Enable Nested Virtualization** check box.

NOTE

Check the note at the bottom of the wizard page before you disable nested virtualization.

Next steps

[Run Hyper-V in a nested VM](#)

Convert a VMware VM to Hyper-V in the VMM fabric

5/20/2019 • 3 minutes to read

This article describes how to convert VMware VMs in the System Center - Virtual Machine Manager (VMM) fabric, to Hyper-V.

There are currently a couple of methods for converting VMWare VMs to Hyper-V:

- **Convert Virtual Machine Wizard:** In the VMM console you can use this wizard. This method has a number of limitations:
 - See [system requirements](#) for supported versions of vSphere (ESXi).
 - You can't convert VMware workstations
 - You can't convert VMs with virtual hard disks connected to an IDE bus
 - Online conversions aren't supported. You need to power off the VMware VMs.
 - Anti-virus apps must be supported.
 - VMware tools must be uninstalled from the guest operating system of the VM.
- **Microsoft Virtual Machine Converter:** This standalone tool converts VMware VMs to Hyper-V hosts or Azure VMs. It also converts physical machines and disks to Hyper-V hosts. IMPORTANT: This tool is in the process of retirement. It won't be available after June 3, 2017. [Learn more](#).

NOTE

Azure Site Recovery currently doesn't have the ability for a direct VMware to Hyper-V conversion. [Read more](#) about up-voting this solution.

Convert using the wizard

1. Click **VMs and Services > Home > Create > Create Virtual Machines > Convert Virtual Machine**.
2. In **Convert Virtual Machine Wizard > Select Source** click **Browse** and in **Select Virtual Machine Source** select the VMware VMs you want to convert.
3. In **Specify Virtual Machine Identity** modify the machine name and description as required.
4. In **Virtual Machine Configuration** specify the number of processor and memory settings.
5. In **Select Host** select a Hyper-V host for placement. In **Select Path** configure the storage location on the host for the VM files. The default VM paths are listed.
6. In **Select Networks** select the logical network, virtual network, and the VLAN as applicable. The list matches whatever is configured on the physical adapters of the host.
7. In **Add Properties** configure settings. In **Summary** review the settings and select **Start the virtual machine after deploying it** if required. Then click **Create** to start the conversion. Verify the VM was converted in **VMs and Services > Home > Show > VMs**.

Convert EFI based VM to Hyper-V generation 2 VM

System Center VMM enables migration of EFI based VMware VMs to Hyper-V. VMware VMs that you migrate to Microsoft Hyper-V platform can now take the advantage of generation 2 features.

As part of VMM 1801 release, the **Convert Virtual Machine** wizard enabled this migration, based on the

firmware type (BIOS or EFI), selects and defaults the Hyper-V VM generation appropriately.

- BIOS-based VMs are migrated to Hyper-V VM Generation 1.
- EFI-based VMs are migrated to Hyper-V VM Generation 2.

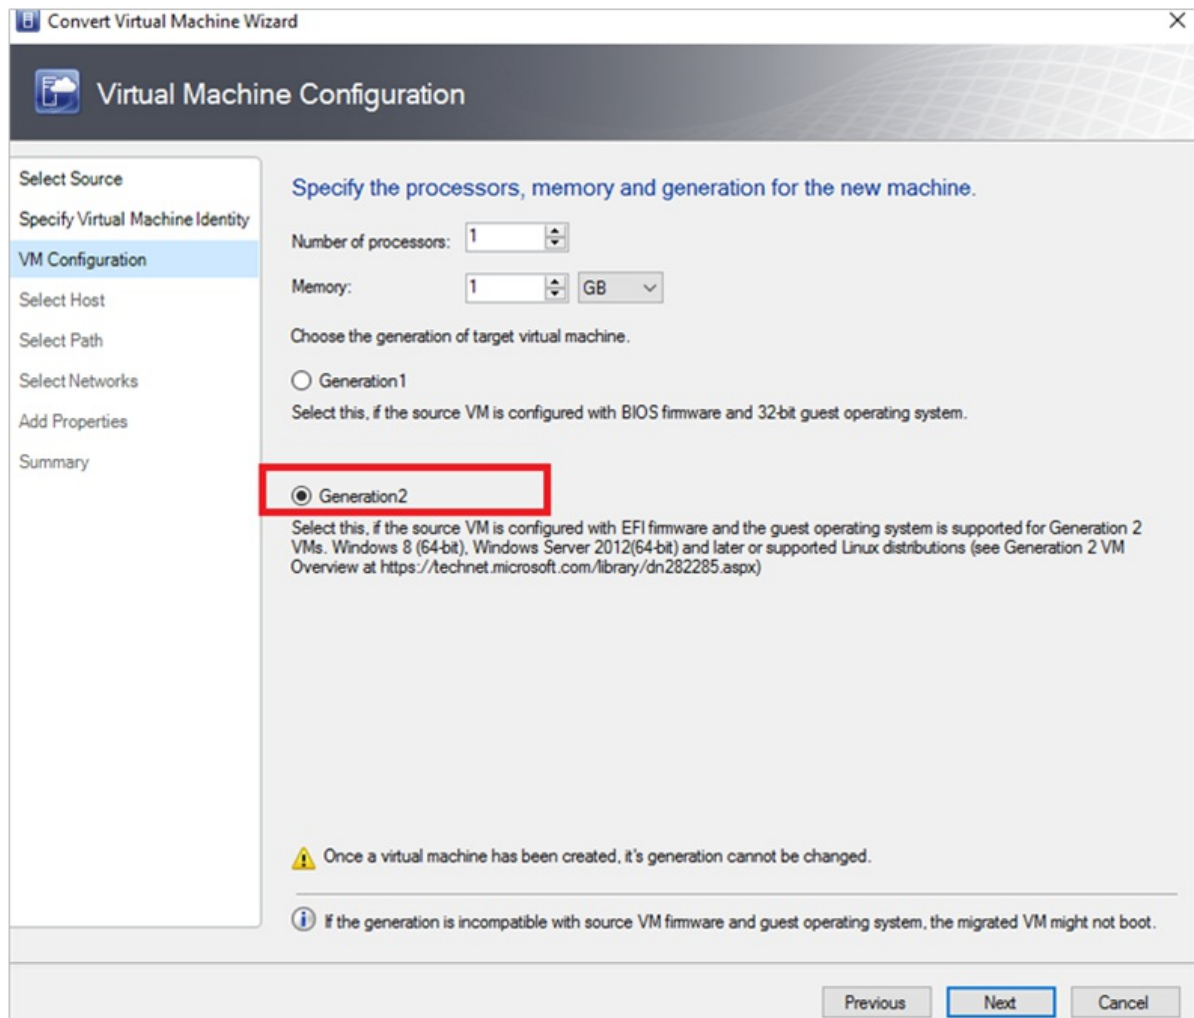
Before you start

Ensure the following prerequisites are met:

1. VMware VMs with firmware type as EFI
2. VMware ESXi Hosts added in System Center VMM

Conversion procedure

1. To convert, follow the [above procedure](#), select **Generation 2** in step 4.



2. Once the VM is converted, you can see the Generation 2 VM as shown in the following image:

Win10Efi Properties

General

Name: Win10Efi

Description:

Computer name:

Cost center:

Tag: (none)

Cloud:

Type: Virtual Machine

Operating system: 64-bit edition of Windows 8.1

Virtual machine guest services: Not Detected

Generation: **Generation 2**

Version: 8.0

Security summary: None

☐ Enable Nested Virtualization

NOTE

- Disk conversion (from "vmdk" to "VHDX/VHD") is enhanced to be ~50% faster than earlier.
- PowerShell commands allow the user to provide the disk type for the target Hyper-V VM, which will enable the VMware thick provisioned disk to be migrated as Hyper-V dynamic disk or vice versa, based upon the requirements.

PowerShell commands

Here are the sample commands:

```
New-SCV2V -VMHost <Host> -VMXPath <string> [-EnableVMNetworkOptimization <bool>] [-EnableMACAddressSpoofing <bool>] [-VMMServer <ServerConnection>] [-LibraryServer <LibraryServer>] [-JobGroup <guid>] [-Trigger] [-VhdType {UnknownType | DynamicallyExpanding | FixedSize}] [-VhdFormat {VHD | VHDX}] [-Description <string>] [-Name <string>] [-Owner <string>] [-UserRole <UserRole>] [-Path <string>] [-StartVM] [-CPUCount <byte>] [-CPURelativeWeight <int>] [-CPUType <ProcessorType>] [-MemoryMB <int>] [-Generation <int>] [-DelayStartSeconds <int>] [-StartAction {NeverAutoTurnOnVM | AlwaysAutoTurnOnVM | TurnOnVMIfRunningWhenVSSStopped}] [-StopAction {SaveVM | TurnOffVM | ShutdownGuestOS}] [-LogicalNetwork <LogicalNetwork>] [-VMNetwork <VMNetwork>] [-NoConnection] [-MACAddress <string>] [-MACAddressType <string>] [-SourceNetworkConnectionID <string>] [-VirtualNetwork <VirtualNetwork>] [-VirtualNetworkAdapter <VirtualNetworkAdapter>] [-VLANEnabled <bool>] [-VLANID <uint16>] [-OverridePatchPath <string>] [-SkipInstallVirtualizationGuestServices] [-NetworkLocation <string>] [-NetworkTag <string>] [-RunAsynchronously] [-PROTipID <guid>] [-JobVariable <string>] [<CommonParameters>]
```

Next steps

[Manage the VM settings](#)

Install an operating system on a VM in the VMM fabric

12/7/2018 • 3 minutes to read

This article describes how to install an operating system on a VM in the System Center - Virtual Machine Manager (VMM) fabric.

After you've deployed a VM in the VMM fabric, you can install an operating system on it. Review the [supported operating systems](#).

You can install an operating system from a DVD, from an ISO image file in the VMM library, or from a network installation.

- For the system CD or ISO image, when you [set up the VM](#) you need a virtual drive to attach to the physical drive or image file.
- To use an ISO image, it must be added to the VMM library.
- For a network installation, you configure a virtual network adapter.

Prepare to install from a system DVD

1. In **Virtual Machines**, right-click the VM > **Properties**.
2. In **Hardware Configuration**, click **DVD** on the **New** toolbar, to add a virtual DVD drive to the IDE bus.
3. Select **Physical CD/DVD drive**, and select the drive on the host.

Prepare to install from an ISO in the VMM library

1. Copy the image file to a share in the Virtual Machine Manager library. The file will be available when the library refreshes. [Learn more](#) about adding file-based resources to the library.
2. In the VM properties > **Hardware Configuration**, click **DVD** on the **New** toolbar, to add a virtual DVD drive to the IDE bus.
3. Select **Known image file**, and select the file from the library resources in the list.

Enable shared ISOs

By default, when you create a VM an ISO attached as a virtual DVD drive is copied into the VM folder. VMM does this so that you can easily migrate VMs from host to host. If you want to share the image from the VMM library, instead of copying it, you can do that as follows:

1. Specify an Active Directory domain account as the VMM service account on the VMM server.
2. Grant the VMM service account read access on the VMM library share that stores the ISO image files. Grant the Hyper-V host machine account read access for the shared ISO location.
3. Configure constrained delegation for each Hyper-V host. This ensures that each host presents delegated credentials for CIFS/SMB, to the VMM server on which the library stores the ISO. To do this, in Active Directory, locate the host machine account and open the account properties. In the **Delegation** tab, select **Select this computer for delegation to specified services only > Use any authentication protocol > Add**. Add the VMM library server that contains the ISO you want to share. In **Add Services**, add **cifs**.
4. Now, configure a virtual machine to share an ISO image. a. In the VM properties > **Hardware Configuration**, in **Capture** mode > select **Existing image file**, and browse to select the ISO image file in the library. b. Select **Share image file instead of copying it**.

NOTE

You must attach the shared ISO image file to the VM after you've created it. You cannot attach the file when you create the VM.

Prepare to install from the network

If the network adapter on the host computer supports network service boots, you can configure a virtual network adapter on the VM to enable this.

1. In **Virtual Machines**, right-click the VM > **Properties**.
2. In **Hardware Configuration**, configure a network connection.
3. On the **New** toolbar, click **Network Adapter** to add a virtual network adapter to the IDE bus.
4. In **Connect to**, select the external virtual network to use for the network service boot. The list contains all virtual networks that are configured on the host.
5. Under **Ethernet (MAC) address**, specify a dynamic or static IP address for the VM.
6. With the VM configured to provide access to the installation medium of choice, you can connect to the virtual machine to install the operating system. By default, VMM uses port 5900 to connect to VMRC. No configuration is required unless a firewall is blocking the port.

Install the operating system on the virtual machine

1. Right-click the VM > **Connect to virtual machine**. Click **Yes** to start the VM.
2. On the Remote Control menu, click **Special Keys** and then click **Send Ctrl+Alt+Delete**.
3. Install the operating system on the VM. The boot disk partition must be the Windows partition.
4. After completing the installation, end your session with the VM and stop the VM in VMM.

Next steps

[Manage VM settings](#).

Configure virtual machine settings in the VMM compute fabric

8/29/2019 • 17 minutes to read

This article describes how to configure performance and availability settings for VMs in the System Center - Virtual Machine Manager (VMM) fabric.

Settings include changing VM properties, and setting up performance options such as quality-of-storage (QoS), availability options, resource throttling, and virtual NUMA.

Add a virtual adapter to a VM

You can add and remove virtual network adapters (vNICs) from VMs that are running. This reduces workload downtime. Note that:

- You add new virtual network adapters by creating or modify a VMM hardware profile.
 - This feature is only available for Generation 2 VMs.
 - By default added virtual network adapters aren't connected to a virtual network. You can configure VMs assigned with the hardware profile to use one or more of the virtual network adapters, after they're deployed on a host.
1. In the virtual machine properties > **Hardware Configuration**, click **Network Adapters**, and select the network adapter you want to add.
 2. You can configure a number of properties for the network adapter, including:
 - **Connected to:** Select what the adapter is connected to.
 - **Not connected:** Select if you don't want to specify a network now.
 - **Internal network:** Select if you want to connect to an isolated internal network that enables communication among VMs on the same host. Virtual machines attached to the internal virtual network cannot communicate with the host, with any other physical computers on the host's LAN, or with the internet.
 - **External network:** Select to specify that a virtual machine created by using this hardware profile will be connected to a physical network adapter on its host. Virtual machines attached to a physical network adapter can communicate with any physical or virtual computer that the host can communicate with and with any resources available on the intranet and over the Internet that the host computer can access.
 - **Ethernet (MAC) address:** computers, a virtual MAC address on virtual machines uniquely identifies each computer on the same subnet. Select one of the following options:
 - **Dynamic.** Select this option if you want to enable a dynamic MAC address for a virtual machine.
 - **Static.** Select this option if you want to specify a static MAC address for a virtual machine. Type a static MAC address in the field provided.

Add a virtual adapter with PowerShell

You can use PowerShell to add a virtual adapter. Here's a sample cmdlet for setting this up:

Example 1

Add a vNIC:

- The first command gets the virtual machine object named VM01, and then stores the object in the \$VM

variable.

- The second command creates a virtual network adapter on VM01.

```
PS C:\> $VM = Get-SCVirtualMachine -Name "VM01"
PS C:\> New-SCVirtualNetworkAdapter -VM $VM -Synthetic
```

Example 2

The following PowerShell commands will remove a vNIC from a running VM. It assumes there's only one vNIC on the VM.

- The first command gets the virtual machine object named VM02, and then stores the object in the \$VM variable.
- The second command gets the virtual network adapter object on VM02, and then stores the object in the \$Adapter variable.
- The last command removes the virtual network adapter object stored in \$Adapter from VM02.

```
PS C:\> $VM = Get-SCVirtualMachine -Name "VM02"
PS C:\> $Adapter = Get-SCVirtualNetworkAdapter -VM $VM
PS C:\> Remove-SCVirtualNetworkAdapter -VirtualNetworkAdapter $Adapter
```

Manage static memory on a running VM

You can modify the memory configuration of a running VM that uses static memory. This feature helps in eliminating workload downtime due to reconfiguration. You can increase or decrease the memory allocation, or switch the virtual machine to dynamic memory. Note that users can already modify dynamic memory for a running VM from VMM and this feature is about modifying the static memory.

Using the following PowerShell examples to modify the static memory setting.

Example 1

Change the static memory for a running virtual machine.

- The first command gets the virtual machine object named VM01, and then stores the object in the \$VM variable.
- The second command changes the memory allocated to VM01 to 1024 MB.

```
PS C:\> $VM = Get-SCVirtualMachine -Name "VM01"
PS C:\> Set-SCVirtualMachine -VM $VM -MemoryMB 1024
```

Example 2

Enable dynamic memory for a running virtual machine.

- The first command gets the virtual machine object named VM02, and then stores the object in the \$VM variable.
- The second command enables dynamic memory, sets the startup memory to 1024 MB, and sets the maximum memory to 2048 MB.

```
PS C:\> $VM = Get-SCVirtualMachine -Name "VM02"
PS C:\> Set-SCVirtualMachine -VM $VM -DynamicMemoryEnabled $True -MemoryMB 1024 -DynamicMemoryMaximumMB 2048
```

Add a servicing window to a VM

You can set up a servicing window for a VM or service so that you can maintain it outside the VMM console. You [set up the window](#), and then assign it to the VM properties.

Create a production checkpoint for a VM

Production checkpoints allow you to easily create "point in time" images of a VM, which can then be restored later.

- Production checkpoints are achieved using backup technology inside the guest to create the checkpoint, instead of using saved state technology.
- On a virtual machine running a Windows operating system production checkpoints are created with the Volume Snapshot Service (VSS).
- Linux virtual machines flush their file system buffers to create a file system consistent checkpoint.
- If you want to create checkpoints using saved state technology you can still choose to use standard checkpoints for your virtual machine.
- You can set one of these checkpoint settings for a VM:
 - **Disabled:** No checkpoint taken.
 - **Production:** Production checkpoints are application consistent snapshots of a virtual machine. Hyper-V leverages the guest VSS provider to create an image of the virtual machine where all of its applications are in a consistent state. The production snapshot does not support the auto-recovery phase during creation. Applying a production checkpoint requires the restored virtual machine to boot from an off-line state just like with a restored backup. This is always more suitable for production environments.
 - **ProductionOnly:** This option is the same as Production with one key difference: With ProductionOnly, if a production checkpoint fails then no checkpoint will be taken. This is different from Production where if a production checkpoint fails, a standard checkpoint will be taken instead.
 - **Standard:** All of the memory state of running applications gets stored so that when you apply the checkpoint the application reverts to the previous state. For many applications this would not be suitable for a production environment. Therefore this type of checkpoint is typically more suitable for development and test environments for some applications.

Set the checkpoint with the following PowerShell command:

```
Set-SCVirtualMachine CheckpointType (Disabled, Production, ProductionOnly, Standard)
```

Configure availability options for clustered VMs

You can configure a number of setting that help high availability and resilience for virtual machines in a cluster:

- **Storage QoS:** You can configure Hyper-V VM hard disks with quality-of-service (QoS) settings, to control bandwidth. You use Hyper-V Manager to do this.
- **Virtual machine priority:** You can configure priority settings for VMs deployed in a host cluster. Based on VM priority, the host cluster starts or places high-priority virtual machines before medium-priority or low-priority virtual machines. This ensures that the high-priority virtual machines are allocated memory and other resources first, for better performance. Also, after a node failure, if the high-priority virtual machines do not have the necessary memory and other resources to start, the lower priority virtual machines will be taken offline to free up resources for the high-priority virtual machines. Virtual machines that are preempted are restarted later in priority order.
- **Preferred and possible owners of virtual machines:** These settings influence the placement of virtual machines on the nodes of the host cluster. By default, there are no preferred owners (there is no preference), and the possible owners include all server nodes on the cluster.

- **Availability sets:** When you place multiple virtual machines in an availability set, VMM will attempt to keep those virtual machines on separate hosts and avoid placing them together on the same host whenever possible. This helps to improve continuity of service.

Configure QoS for a VM

1. Open **Hyper-V Manager**, and click **Action > Settings**.
2. In **SCSI Controller**, click **Hard Drive**
3. In **Advanced Features**, click **Enable Quality of Service management**.
4. Specify minimum and maximum IOPS values.

Configure priority

1. Configure a virtual machine or virtual machine template by using one of the following options:
 - To configure a deployed virtual machine, in **VMs and Services**, navigate to the host on which the virtual machine is deployed. Right-click the virtual machine > **Properties**.
 - To configure a stored virtual machine, in the **Library**, navigate to the library server on which the virtual machine is stored. Right-click the virtual machine > **Properties**.
 - You can also set up priority while you're configuring a VM, on the **Configure Hardware** page. To configure a virtual machine template, in **Library > Templates**, click **VM Templates**. Right-click the virtual machine template, > **Properties**.
2. In **Hardware Configuration** or **Configure Hardware**, scroll down to **Advanced**, and click **Availability**. Make sure **Make this virtual machine highly available** is checked. On a deployed virtual machine this setting cannot be changed, because it depends on whether the virtual machine is deployed on a host cluster.
3. In **Virtual machine priority**, select a priority of High, Medium, or Low for the VM. If you want the virtual machine to always require a manual start and never preempt other virtual machines, select **Do not restart automatically**.

Configure preferred owners

1. In **VMs and Services**, navigate to the host on which the virtual machine is deployed. Right-click the virtual machine > **Properties**.
2. Click **Settings** and configure the options:
 - To control which nodes (servers) in the cluster will own the virtual machine most of the time, configure the preferred owners list.
 - To prevent a virtual machine from being owned by a particular node, configure the possible owners list, omitting only the nodes that should never own the virtual machine.

Configure availability sets

You can configure availability sets for standalone VMs in a cluster, or in availability sets in a service template, to specify how VMs created with the template should be placed on hosts.

1. Configure a virtual machine or virtual machine template by using one of the following options:
 - To configure a deployed virtual machine, in **VMs and Services**, navigate to the host on which the virtual machine is deployed. Right-click the virtual machine > **Properties**.
 - To configure a stored virtual machine, in the **Library**, navigate to the library server on which the virtual machine is stored. Right-click the virtual machine > **Properties**.
 - You can also set up priority while you're configuring a VM, on the **Configure Hardware** page. To configure a virtual machine template, in **Library > Templates**, click **VM Templates**. Right-click the virtual machine template, > **Properties**.
2. On the **Hardware Configuration** tab, scroll down to **Advanced** and under it, click **Availability**.

3. Confirm that **Make this virtual machine highly available** has the intended setting. (On a deployed virtual machine, the setting cannot be changed, because it depends on whether the virtual machine is deployed on a host cluster.)
4. Under **Availability sets**, click **Manage availability sets**.
5. Click the name of an availability set, and use the controls to add or remove the set. Repeat this action until all of the intended availability sets appear in the **Assigned properties** list. To create a new availability set, click the **Create** button, provide a name for the set, and then click **OK**.
6. To verify the setting for a deployed virtual machine, in the listing for the virtual machine, view the name under **Availability Set Name**.

For virtual machines that have been deployed on a host cluster, another way to configure this setting is to use Windows PowerShell commands for failover clustering. In this context, the setting appears in the [Get-ClusterGroup](#) as **AntiAffinityClassNames**.

Configure resource throttling

VMM includes resource throttling features such as processor (CPU) and memory throttling, to control resource allocation and help virtual machines to run more effectively.

- **Processor throttling:** You can set the weight of a virtual processor to provide the processor with a larger or smaller share of CPU cycles. The properties ensure that VMs can be prioritized or deprioritized when CPU resources are overcommitted. For highly intensive workloads, more virtual processors can be added, especially when a physical CPU is close to its upper limit.
 - **High, Normal, Low, Custom:** Specifies how the CPU is distributed when contention occurs. Higher priority virtual machines will be allocated CPU first.
 - **Reserve CPU cycles (%):** Specifies the percentage of CPU resources that are associated with one logical processor that should be reserved for the virtual machine. This is useful when a virtual machine runs applications that are particularly CPU-intensive and you want to ensure a minimal level of CPU resources. A zero setting indicates that no specific CPU percentage is reserved for the virtual machine.
 - **Limit CPU cycles (%)** Specifies that the virtual machine should not consume more than the indicated percentage of one logical processor.
- **Memory throttling and weight:** Memory throttling helps to prioritize or deprioritize access to memory resources in scenarios where memory resources are constrained. When memory usage on a host is high, then the virtual machines with a higher memory priority are allocated memory resources before the virtual machines with a lower priority. If you specify a lower priority, it might prevent a virtual machine from starting when other virtual machines are running and the available memory is low. You can set the memory priority settings and thresholds as follows:
 - **Static:** The amount of static memory that is assigned to a specific virtual machine.
 - **Dynamic:** Dynamic memory settings include:
 - **Start-up memory:** The amount of memory that is allocated to the virtual machine when it starts up. It should at least be set to the minimum amount of memory that is required to run the operating system and applications on the virtual machine. Dynamic memory will adjust the memory amount as required.
 - **Minimum memory:** The minimum amount of memory that is required for the virtual machine. It allows an idle machine to scale back the memory consumption below the start-up memory requirement. The available memory can then be used by other virtual machines.
 - **Maximum memory:** The memory limit that is allocated to the virtual machine. The default value is 1 TB.

- **Memory Buffer Percentage:** Dynamic memory adds memory to a virtual machine as required, but there is a chance that an application might demand memory more quickly than dynamic memory allocates it. The memory buffer percentage specifies the amount of available memory that will be assigned to the virtual machine if needed. The percentage is based on the amount of memory that is actually needed by the applications and services that run on the virtual machine. It is expressed as a percentage because it changes depending on the virtual machine requirements.
 - The percentage is calculated as follows: $\text{Amount of memory buffer} = \frac{\text{memory needed by the virtual machine}}{(\text{memory buffer value}/100)}$.
 - For example, if the memory that is committed to the virtual machine is 1000 MB and the buffer is 20%, then an additional buffer of 20% (200 MB) will be allocated for a total of 1200 MB of physical memory allocated to the virtual machine.
- **Memory weight:** The priority that is allocated to a virtual machine when the memory resources are in full use. If you set a high priority value, it will prioritize a virtual machine when the memory resources are allocated. If you set a low priority, a virtual machine might be unable to start if memory resources are insufficient.

Configure processor throttling

1. In the virtual machine > **Properties** > **Advanced**, click **CPU Priority**.
2. Select a priority value for the virtual machine. These values specify how the CPU resources are balanced between virtual machine, and correspond to the relative weight value in Hyper-V:
 - High—Relative weight value of 200
 - Normal—Relative weight value of 100
 - Low—Relative weight value of 50
 - Custom—Relative weight values that are supported are between 1 and 10000
3. In **Reserve CPU cycles (%)**, specify the percentage of the CPU resources on one logical processor that should be reserved for a virtual machine. This is useful when a virtual machine runs applications that are particularly CPU-intensive, and you want to ensure a minimal level of CPU resources. A zero setting indicates that no specific CPU percentage is reserved.
4. In **Limit CPU cycles (%)**, specify the maximum percentage of the CPU resources on one logical processor that the virtual machine should consume. The virtual machine will not be allocated more than this percentage.

Configure memory throttling

1. In the virtual machine > **Properties** > **General**, click **Memory**.
2. Select **Static** to specify that a fixed amount of memory should be assigned to a virtual machine.
3. Select **Dynamic** to specify the dynamic memory settings for a virtual machine, as follows:
 - In **Startup memory**, specify the amount of memory that is allocated to the virtual machine when it starts up. The memory value should be set at least to the minimum amount of memory that is required for the virtual machine operating system and applications to run.
 - In **Minimum memory**, specify an amount of memory that allows an idle virtual machine to scale back the memory consumption below the startup memory requirement. This makes more memory available for use by other virtual machines.
 - In **Maximum memory**, specify the maximum amount of memory that is allocated to a virtual machine. The default setting is 1 TB.
 - In **Memory buffer percentage**, specify the amount of available memory that will be assigned to a virtual machine if the need arises. The percentage should be based on the amount of memory that is actually needed by the applications and services that run on the virtual machine. The memory buffer

percentage should be calculated as follows: Amount of memory buffer = memory that is needed by the virtual machine/ (memory buffer value/100). For example, if the memory that is committed to the virtual machine is 1000 MB and the buffer is 20%, then an additional buffer of 20% (200 MB) will be allocated for a total of 1200 MB of physical memory allocated to the virtual machine.

Configure virtual NUMA

You configure, deploy, and manage virtual Non-Uniform Memory Access (NUMA) in VMM. Virtual NUMA has the following properties:

- NUMA is a memory architecture that is used in multiprocessor systems, where the time that is required for a processor to access memory depends on the location of the memory relative to the processor. On a NUMA system, a processor can access the local memory (the memory that is directly attached to the processor) faster than the non-local memory (the memory that is attached to another processor). NUMA attempts to close the gap between the speed of processors and the memory that they use. To do so, NUMA provides separate memory on a per-processor basis, thus This helps to avoid the performance degradation that occurs when multiple processors try to access the same memory. Each block of dedicated memory is known as a NUMA node.
- Virtual NUMA enables the deployment of larger and more mission-critical workloads that can be run without significant performance degradation in a virtualized environment, when compared to running non-virtualized computers with physical NUMA hardware. When a new virtual machine is created, by default Hyper-V uses values for the guest settings that are in sync with the Hyper-V host NUMA topology. For example, if a host has 16 cores and 64 GB divided evenly between two NUMA nodes with two NUMA nodes per physical processor socket, then a virtual machine that is created on the host with 16 virtual processors will have the maximum number of processors per node setting set to eight, maximum nodes per socket set to two, and maximum memory per node set to 32 GB.
- NUMA spanning can be enabled or disabled. With spanning enabled, individual virtual NUMA nodes can allocate non-local memory, and an administrator can deploy a virtual machine that has more virtual processors per virtual NUMA node than the number of processors that are available on the underlying hardware NUMA node on the Hyper-V host. NUMA spanning for a virtual machine does incur a performance cost because virtual machines access memory on non-local NUMA nodes.

Set up virtual NUMA for VMs as follows:

1. In the virtual machine > **Properties** > **Advanced**, click **Virtual NUMA**.
2. In **Maximum processors per virtual NUMA node**, specify the maximum number of virtual processors that belong to the same virtual machine and that can be used concurrently on a virtual NUMA node. Configure this setting to ensure maximum bandwidth. different NUMA virtual machines to use different NUMA nodes. The minimum limit is 1 and the maximum is 32.
3. In **Maximum memory per virtual NUMA node (MB)**, specify the maximum amount of memory (MB) that can be allocated to a single virtual NUMA node. The minimum limit is 8 MB and the maximum is 256 GB.
4. In **Maximum virtual NUMA nodes per socket**, specify the maximum number of virtual NUMA nodes that are allowed on a single socket. The minimum number is 1 and the maximum is 64.
5. To enable spanning, click **Allow virtual machine to span hardware NUMA nodes**.

Set up dynamic and power optimization in VMM

5/30/2019 • 9 minutes to read

Read this article to learn about enabling dynamic optimization (DO) and power optimization for virtual machines (VMs) in the System Center - Virtual Machine Manager (VMM). The article includes features overview, instructions for setting up BMC for power optimization, and describes how to enable and run these features.

NOTE

VMM 2019 and later supports dynamic optimization for Compute and Storage. Versions prior to VMM 2019 support DO for compute only, use the following procedures as applicable for the version of VMM you are using.

- **Dynamic optimization:** Using dynamic optimization, VMM performs live migration of VMs and VHDs within a host cluster. The migration is based on the settings you specify, to improve load balancing among hosts and cluster shared storage (Cluster shared volumes (CSV), file shares), and to correct the placement issues for VMs.
 - **Compute Dynamic optimization** (Optimization of hosts) can be performed on hosts in a cluster to optimize host performance by migrating VMs across hosts. Host performance thresholds you can set are: CPU and Memory.
 - **Storage Dynamic Optimization** (Optimization of disk space, applicable for VMM 2019 and later) can be performed on cluster shared storage (CSV, file shares) to optimize storage space availability by migrating Virtual Hard Disks (VHD) across shared storage. You can set free storage space threshold on cluster shared storage.
- **Power optimization:** Power optimization is a feature of dynamic optimization that saves energy by turning off hosts that aren't needed to meet resource requirements within a cluster, and turns them back on when they're needed.

VMM supports compute dynamic optimization (both compute and storage in VMM 2019 and later) and power optimization on Hyper-V host clusters. Compute dynamic optimization and power optimization is also supported on VMware host clusters in the VMM fabric that support live migration.

Before you start

Note the following information before you start using the DO.

Dynamic optimization

- Dynamic optimization and power optimization can be configured on host clusters that support live migration.
- Dynamic optimization can be configured on a host group, to migrate virtual machines and virtual hard disks (VHDs) within host clusters with a specified frequency and aggressiveness. VM aggressiveness determines the amount of load imbalance that is required to initiate a migration during dynamic optimization.
- Disk space aggressiveness determines the amount of free storage space below disk space threshold that is required to migrate VHDs to other cluster shared storage during dynamic optimization.
- By default, virtual machines are migrated every 10 minutes with medium aggressiveness, if automatic migration is enabled. When configuring frequency and aggressiveness for dynamic optimization, an administrator should factor in the resource cost of additional migrations against the advantages of balancing

load among hosts/shared storage in a host cluster. By default, a host group inherits Dynamic Optimization settings from its parent host group.

- If you set up dynamic optimization on a host group without a cluster it will have no effect.
- Dynamic optimization can be set up for clusters with two or more nodes. Storage dynamic optimization will need two or more shared storage files/volumes to be present in the cluster. If a host group contains stand-alone hosts or host clusters that do not support live migration, dynamic optimization isn't performed on those hosts. Any hosts that are in maintenance mode also are excluded from dynamic optimization. In addition, VMM only migrates highly available virtual machines that use shared storage. If a host cluster contains virtual machines that are not highly available, those virtual machines are not migrated during Dynamic Optimization.
- On-demand dynamic optimization is also available for individual host clusters by using the Optimize Hosts/ Optimize Disk space action in the VMs and Services workspace. It can be performed without configuring dynamic optimization on host groups. After dynamic optimization is requested for a host cluster, VMM lists the virtual machines/VHDs that will be migrated, for the administrator's approval. Optimize Hosts performs VM load balancing across hosts in a cluster, while Optimize disk space migrates VHDs across Shared storage in a cluster.

Node fairness

Node fairness is a new feature in Windows Server 2016. It identifies cluster nodes with light loads, and distributes VMs to those node to balance load. This is similar to VMM's dynamic optimization. To avoid potential performance issues, dynamic optimization and node fairness shouldn't work together. To ensure this doesn't happen VMM disables node fairness in all clusters in a host group for which dynamic optimization is set to automatic. If you enable node fairness outside the VMM console, VMM will turn it off the next time that dynamic optimization refreshes. If you do want to use node fairness, disable dynamic optimization, and then manually enable node fairness.

Power optimization

- For power optimization, the computers must have a baseboard management controller (BMC) that enables out-of-band management.
- Power optimization ensures that a cluster maintains a quorum if an active node fails. For clusters created outside VMM and added to VMM, Power Optimization requires more than four nodes. For each additional one or two nodes in a cluster, one node can be powered down. For instance:
 - One node can be powered down for a cluster of five or six nodes.
 - Two nodes can be powered down for a cluster of seven or eight nodes.
 - Three nodes can be powered down for a cluster of nine or ten nodes.
- When VMM creates a cluster, it creates a quorum disk and uses that disk as part of the quorum model. For clusters created by VMM, Power Optimization can be set up for clusters of more than three nodes. This means that the number of nodes that can be powered down is as follows:
 - One node can be powered down for a cluster of four or five nodes.
 - Two nodes can be powered down for a cluster of six or seven nodes.
 - Three nodes can be powered down for a cluster of eight or nine nodes.

Configure BMC

For hosts with BMC that supports IMPI 1.5/2.0, DCMI 1.0 or SMASH 1.0 over WS-Management you can figure BMC settings as follows:

1. Create a Run As account with permissions to access the BMC on a host.
2. Click **Fabric > Servers > All Hosts > host > Properties > Hardware > Advanced > BMC Setting**.
3. To enable VMM management, select **This physical machine is configured for OOB management**.
4. In **This computer supports the specified OOB power management configuration provider**, click the

supported management protocol. Type in the IP address of the BMC, and accept the default port offered by VMM. Select the Run As account and click **OK**.

Enable dynamic and power optimization for a host group

1. Click **Fabric > Servers > All Hosts**, and select the host group that you want to configure.
2. With the host group selected, click **Folder > Properties group > Properties**.
3. In the host group properties, click **Dynamic Optimization**.
4. In **Specify dynamic optimization settings**, clear the **Use Dynamic Optimization settings from the parent host group** check box.
5. In **Aggressiveness**, select **High**, **Medium**, or **Low**.

VM aggressiveness determines the amount of load imbalance that is required to initiate a migration during dynamic optimization.

Disk space aggressiveness determines the amount of free storage space below disk space threshold that is required to migrate VHDs to other cluster shared storage during dynamic optimization.

When you configure frequency and aggressiveness for dynamic optimization, you should try to balance the resource cost of additional migrations against the advantages of balancing load among hosts in a host cluster. Initially, you might accept the default value of **Medium**. After you observe the effects of dynamic optimization in your environment, you can increase the aggressiveness.

6. To help conserve energy by having VMM turn off hosts when they are not needed and turn them on again when they are needed, configure power optimization for the host group. Power optimization is only available when virtual machines are being migrated automatically to balance load.
7. To periodically run dynamic optimization on qualifying host clusters in the host group, enter the following settings:
 - a. Select the **Automatically migrate virtual machines to balance load** check box to balance free storage space across shared storage.
 - b. In **Frequency**, specify how often to run dynamic Optimization. You can enter any value between 10 minutes and 1440 minutes (24 hours).
8. Set thresholds for each of the compute and storage (applicable for VMM 2019) resources listed. To change the units of the resources go to **Host group > Properties > Host Reserves** and choose the unit from the drop-down menu.
9. To turn on power optimization on the host group, select the **Enable power optimization** checkbox. Click **OK** again to save your changes.

NOTE

If there is a mismatch of disk space warning levels between host groups having the same file share, it can result in multiple migrations to and from that file share and may impact storage DO performance. We recommended you not to do a file share across different host groups where storage dynamic optimization is enabled.

Configure power optimization settings

1. In the **Fabric** navigate to the host group and open **Properties**.
2. Click **Dynamic Optimization > Specify dynamic optimization settings > Settings**.
3. In **Customize Power Optimization Schedule**, change the settings for any of these resources: CPU, memory,

disk I/O , or network I/O.

4. Under **Schedule**, select the hours when you want power optimization to be performed. Click a box to turn power optimization on or off for that hour. VMM applies the schedule according to the host time zone.

Run dynamic optimization on-demand in a host cluster

You can run dynamic optimization on demand on a host cluster. To do this dynamic optimization doesn't need to be configured on the parent host group.

1. Open **Fabric > Servers > Host Groups**, and navigate the host cluster.
2. To perform compute resource load balancing, click **Optimize hosts**. To perform storage load balancing across cluster shared storage, click **Optimize disks**.

To Optimize hosts: VMM performs a dynamic optimization review to determine whether VHDs can be migrated to improve load balancing in the host cluster. If migration of VMs can improve load balancing, VMM displays a list of VMs that are recommended for migration, with the current and target hosts indicated. The list excludes any hosts that are in maintenance mode in VMM and any virtual machines that are not highly available.

To Optimize Disk Space: VMM performs a dynamic optimization review to determine whether VHDs can be migrated to meet the free storage space threshold (disk space) while considering aggressiveness set in the Dynamic Optimization page. Dynamic Optimization will only be triggered when any cluster shared storage violates the disk space threshold set. If migration of VHDs can help free the storage space threshold in shared storage in the cluster, VMM displays a list of VHDs that are recommended for migration, with the current and target storage space indicated. VHDs will only migrate to another shared storage with the same storage classification.

3. Click **Migrate**.

NOTE

If VHDs are migrated between one storage type to another (Example: from a CSV to NAS file share), the storage migration will be slow. If the storage optimization does not return a list of VHDs to migrate even when the threshold and aggressiveness criteria are met: - Check the HostVolumeID using Get-SCStorageVolume Cmdlet. If the HostVolumeID returns Null for the volume, refresh the VM and perform Storage Dynamic Optimization again. - Check the DiskSpacePlacementLevel of the host group using the Get-SCHostReserver cmdlet. Set the DiskSpacePlacementLevel value equal to the value of Disk Space set in Host Reserve settings in the Dynamic Optimization wizard.

Power on/off a computer in VMM

1. Click **Fabric > Servers > All Hosts > host name**.
2. On the **Host** tab, in the **Host** group click **Power On** or **Power Off**. You can view information about power on and off events in the BMC logs (click on **Hardware > Advanced > BMC Logs**).

Next steps

Learn about [provisioning VMs](#).

Create VM templates using VMM and Windows Azure Pack

8/29/2019 • 5 minutes to read

As a hosting provider, you can use System Center - Virtual Machine Manager (VMM) in combination with Windows Azure Pack to increase the features that you offer to tenants. To help tenants create virtual machines with specific operating systems and applications already installed, you can create VM role templates. Tenants can use these templates to create VMs on-premises, and in service-provider hosting environments.

Before you begin

- Learn about [Windows Azure Pack](#)
- Learn about how VMM manages resources in the [VMM library](#)

Download Azure Pack Gallery resources

Azure Pack Gallery Resources provide tenant offerings using standard and reusable components. The VM role gallery enables you to deploy VMs. Usually you need two packages to deploy a VM role:

- A resource definition package (*.resdefpkg) to deploy the VM. This template describes information asked to self-portal tenants about VM size, name etc. These parameters are used by VMM to deploy the role.
- A resource extension package (*.resextpkg) to install apps on a VM.

Each resource contains a Readme file that explains how to prepare your environment.

1. Download the [Microsoft Web Platform Installer](#). You use this to download the gallery resources.
2. Start the installer, click the **Options** link at the bottom, and type the link <https://www.microsoft.com/web/webpi/partners/servicemodels.xml> into **Custom Feeds**. Click **Add Feed > OK**. A **Service Models** tab should appear.
3. In the **Service Models** tab, select **Gallery Resources**.
4. Select the resource you want, and click **Add > Install**.
5. Accept the license terms if appropriate, and then click **Continue**. A window will be opened for the folder on the local computer where the resource is downloaded.
6. The Readme file that accompanies the resource will specify whether additional software is needed. Follow the instructions as needed.

Import resource extensions to the library

For gallery resources that use Resource Extensions, you need to import the extensions into the VMM library, using PowerShell. This example shows you how to install a Resource Extension Pack (MyVMRole.resextpkg), with a library shared named MSSCVMMLibrary,

1. At the command line, run the following:

```
$libraryShare = Get-SCLibraryShare | Where-Object {$_.Name -eq 'MSSCVMMLibrary'}
$resextpkg = $Env:SystemDrive + "\GalleryResources\My-VMRole-Pkg\MyVMRole.resextpkg"
Import-CloudResourceExtension -ResourceExtensionPath $resextpkg -SharePath $libraryShare -
AllowUnencryptedTransfer
```

2. Verify that the import completed with the following cmdlet:

```
Get-CloudResourceExtension
```

Create and prepare a virtual hard disk

To create a template, you need a virtual hard disk file that contains an operating system that's been prepared for deployment, using Sysprep.

Specify the operating system

1. When you create the hard disk, specify the required operating system in the Operating System property. For data disks this value must be set to None.
2. Set the property using the VMM console, or PowerShell. The following example shows how to set the virtual hard disk MyVirtualHardDisk to run Windows Server 2012 Datacenter. Your cmdlet with replace the disk name with one of the values in the Readme file.

```
$myVHD = Get-SCVirtualHardDisk | where {$_.Name -eq 'MyVirtualHardDisk.vhd'}
$WS2012Datacenter = Get-SCOperatingSystem | where { $_.name -eq '64-bit edition of Windows Server 2012 Datacenter' }
Set-scvirtualharddisk -virtualharddisk $myVHD -OperatingSystem $WS2012Datacenter
```

Specify the name and release

1. The Familyname and Release properties of the virtual hard disk must be set in order for the Windows Azure Pack portal to display the virtual hard disk as an available disk for this Gallery Resource. These values are shown in the portal drop-down list.
 - The Familyname property values should indicate the contents of the virtual hard disk, including the Windows Server release and edition. The Readme file of the Gallery Resource should include appropriate Familyname values. Release property values must conform to the Windows Azure versioning scheme of n.n.n.n. Examples include 1.0.0.0 and 1.0.0.1.
2. Set the property using the VMM console, or PowerShell. The following example shows how to set the virtual hard disk MyVirtualHardDisk to a Familyname "Windows Server 2012 Datacenter", and the Release property to "1.0.0.0". Your cmdlet with replace values with one of the values in the Readme file.

```
$myVHD = Get-SCVirtualHardDisk | where {$_.Name -eq 'MyVirtualHardDisk.vhd'}
$familyName = "Windows Server 2012 DataCenter"
$release = "1.0.0.0"
Set-scvirtualharddisk -virtualharddisk $myVHD -FamilyName $familyName -Release $release
```

Specify tags

Virtual Machine Role gallery items specify tags that must be included on an operating system disk for it to be available as an option when a user provisions the virtual machine. The Readme file of the gallery resource should include the tags that it requires.

1. When you create the hard disk, specify the required tags using PowerShell.
2. The following example shows how to set the tag "WindowsServer2012R1" for the virtual hard disk MyVirtualHardDisk. Your cmdlet use the values in the Readme file.

```
$myVHD = Get-SCVirtualHardDisk | where {$_.Name -eq 'MyVirtualHardDisk.vhd'}
$tags = $myVHD.Tag
if ( $tags -cnotcontains "WindowsServer2012R1" ) { $tags += @"WindowsServer2012R1" }
Set-scvirtualharddisk -virtualharddisk $myVHD -Tag $tags
```

Add the virtual disk to the library

After you create the virtual hard disk, you must add it to the VMM library with the values you specified. [Learn more](#) about adding file-based resources to the library.

Import the Resource Definition packages

After the Resource Extension and virtual hard disk are in the VMM library, you can import the Resource Definition package, and publish the gallery item using the Service Administrator Portal in Windows Azure Pack. After these steps, the gallery item will be available to the tenant.

- The Gallery Resource will include one or more Resource Definition Package files.
 - If it includes more than one, then the ReadMe file will specify the different configuration that each one will provide.
1. Open the Service Admin Portal.
 2. Navigate to the **VM Clouds** workspace.
 3. Click the **Gallery** tab > **Import**.
 4. Select and import the Resource Definition Package file for the gallery item. This will be the file with the extension resdefpkg.

The gallery item should now be listed on the **Gallery** tab.

Publish the item and add to a plan

1. On the **Gallery** tab, select the version of the gallery item that you just imported.
2. Click the arrow next to the gallery item name. Verify the details of the gallery item. Navigate back, and click **Make Public**.
3. Select the **Plans** workspace, and select the plan to which you want to add the gallery item.
4. Select the **Virtual Machine Clouds** service. Scroll to the **Gallery** section, and click **Add Gallery Items**. 5 Select the gallery items that you imported, and then click **Save**

The Virtual Machine Role is now available to the tenant as part of the selected plan.

Next steps

Learn about [provisioning VMs](#).

Migrate virtual machines in the VMM fabric

12/20/2018 • 7 minutes to read

This article provides an overview of migrating virtual machines in the System Center - Virtual Machine Manager (VMM) fabric.

You can migrate virtual machines and storage managed in the VMM fabric. VMM automatically selects the type of transfer that will be used for migration. When you perform a migration in the VMM console using the Migrate VM Wizard, the migration type that will be used is displayed in the **Transfer Type** column. The types of migrated supported are summarized in the following table.

TYPE	USE	DETAILS
Network migration	Performs a network copy of the virtual machine data using BITS.	This is the slowest type of migration. The amount of downtime is in direct proportion to the size of the data transfer.
Quick migration	Also known as cluster transfer, and can be used to migrate a highly available virtual machine. It leverages Windows Failover Cluster to migrate virtual machines between cluster nodes.	<p>The running state of the virtual machine is saved to disk (the virtual machine is hibernated), the disk is failed over to the other cluster node, and then the saved state is loaded to wake up the virtual machine.</p> <p>Downtime is minimal because quick migration takes a snapshot of the virtual machine and transfers data without requiring the virtual machine to be turned off.</p>
Quick storage migration	Use to move VM storage from one location to another. For example, you can move the storage for a virtual machine from a Fibre Channel SAN to an iSCSI SAN.	<p>The virtual disks of a running virtual machine can be migrated independent of storage protocols (SCSI, Fibre Channel) or storage types (local, DAS, SAN).</p> <p>Downtime is minimal because quick storage migration takes a snapshot of the virtual machine and transfers data without requiring the virtual machine to be turned off.</p>

TYPE	USE	DETAILS
SAN migration	Uses SAN transfer to migrate virtual machines, and highly available virtual machines, in and out of a cluster. It can be used when both the source and destination hosts have access to the same storage infrastructure (LUN), and the storage can be transferred from one host to another.	<p>For SAN migration, the files for a virtual machine are not copied from one server to another and thus downtime is minimized. SAN migration can be used to copy a virtual machine from one host to another, or copying a virtual machine to or from the library.</p> <p>When you migrate a virtual machine into a cluster by using a SAN transfer, VMM checks that each node in the cluster can see the LUN, and automatically creates a cluster disk resource for the LUN.</p> <p>To migrate a virtual machine out of a cluster, the virtual machine must be on a dedicated LUN that is not using CSV.</p> <p>These SAN infrastructures are supported for migration: Fiber Channel; iSCSI SANs; N_Port ID Virtualization (NPID).</p>
Live migration	Moves a virtual machine running as part of a failover cluster from one cluster to another.	No noticeable downtime for users or network applications.

Live migration

Using live migration provides a number of benefits:

- **Increased flexibility:** Live migration features can help to simplify the movement of virtual machines across hosts and clusters. Therefore, it becomes easier to manage a dynamic datacenter.
- **Ease-of-maintenance:** Live migration alleviates the need to take standalone hosts and cluster hosts offline for maintenance and migration purposes, which helps to avoid downtime. With the ability to perform concurrent migrations and maintenance, migration timeframes can become shorter, depending on the time that is required to perform the live migration. In addition, the planning process for Hyper-V mobility is simplified.
- **Better hardware utilization:** The distribution of virtual machines can be optimized across the infrastructure. Virtual machines and storage can be moved to stand-alone servers and clusters with spare capacity, without interrupting availability. Power consumption is reduced as virtual machines can be moved across hosts, and then hosts can be powered down to save energy.
- **Failover clustering features:** VMM takes advantage of failover clustering features that were introduced in Windows Server 2012. These features include additional APIs to migrate virtual machines across cluster nodes, and improved attach/detach functionality that enables migration of virtual machines in and out of failover clusters without downtime. VMM supports the following types of live migration:

Live migration support

VMM supports the following types of live migration:

- **Live migration of standalone machines:** You can run live migration between two standalone machines that aren't in a cluster.
- **Live migration within a cluster:** You can run a live migration between nodes in the same cluster.
- **Live migration between nodes in different clusters:** You can migrate between nodes in different clusters.
- **Live migration of VM storage:** You can migrate storage to update the physical storage available in Hyper-V,

or to mitigate bottlenecks in storage performance. You can also use storage migration to move, service, or upgrade storage resources, or for migration of a standalone or cluster virtual machine. Storage can be added to either a standalone computer or to a Hyper-V cluster. VMs can be moved to the new storage while they continue to run.

- **Live VSM:** You can use live system migration (live VSM) to migrate virtual machines and their storage together, in a single action.
- **Concurrent live migration:** You can perform multiple concurrent live migrations of virtual machines and storage. The concurrent limit can be configured manually. Any concurrent live migrations in excess of the limit will be queued.

VMM inspects and validates the configuration settings of a destination host before migration from a source host begins.

Live VM migration support matrix

SOURCE	DESTINATION: STANDALONE	DESTINATION: CLUSTER
Standalone	Supported	Supported
Cluster	Supported	Supported Source and destination can be in the same or different clusters.

Live storage migration support matrix

SOURCE	DESTINATION: LOCAL DISK (STANDALONE)	DESTINATION: SMB 3.0 SHARE (STANDALONE/CLUSTER)	DESTINATION: CSV (CLUSTER)
Local disk	Supported	Supported. The virtual machine will be promoted to high availability.	Not supported.
SMB 3.0 share	Supported. In a cluster the VM will be demoted, and won't be highly available after migration.	Supported	Supported
Cluster	Supported In a cluster the VM will be demoted, and won't be highly available after migration.	Supported The SMB share must be available from the destination cluster node.	Supported The CSV must be available from the destination cluster node.

Live migration limitations

- Live migration requires two or more servers that run Hyper-V, that support hardware virtualization, and use processors from the same manufacturer, such as all AMD processors or all Intel processors, for example.
- Live migration is supported starting with hosts running Windows Server 2012.
- Virtual machines must be configured to use virtual hard disks or virtual Fibre Channel disks, not physical disks.
- For live migration network traffic you should use a private network.
- Source and destination servers must belong to the same Active Directory domain, or to different trusted

domains.

- If the source or destination virtual machine VHD has a base disk, the base disk must be in a share that is accessible (registered) from the destination host. Generally, live migration does not move the base disk.
- Migration between clusters is only supported on hosts running in failover clusters. Cluster Shared Volume (CSV) storage should be enabled in the cluster.
- Live migration of a virtual machine does not migrate virtual machine storage, specifically meaning the location that stores the virtual machine images (VHD, ISO, VFD files). To handle storage requirements, you can use one of the following options:
 - Configure the virtual machine so that the storage files are available on a file share that is accessible by both the source and destination host of the migration.
 - Run a combined live virtual machine and storage migration (live VSM) in a single action.
 - Run a separate storage migration.
- If the source and destination hosts use shared storage, note the following:
 - All files that comprise a virtual machine, such as virtual hard disks, snapshots, and configuration, should be stored on an SMB share.
 - Permissions on the SMB share should be configured to grant access to the computer accounts of all servers that run Hyper-V.
- A storage migration moves virtual machine images (VHD, ISO, and VFD files), snapshot configurations, and data (saved state files).
- Storage migration is per virtual machine.
- Storage migration does not move base (parent) disks, with the exception of snapshot disks.

Live VSM

Live VMS migrates a VM and its machine storage in a single action.

- To use live VSM, the virtual machine LUN must be masked from the destination host.
- Live VSM is supported between two standalone hosts that run Hyper-V, starting with Windows Server 2012. The transfer can occur between local disks or SMB 3.0 file shares.
- Live VSM is supported between two host clusters that run Hyper-V, starting with Windows Server 2012. The virtual machine can be transferred to either a CSV or SMB 3.0 file share on the destination host cluster.

Next steps

- [Migrate a virtual machine](#)
- [Migrate storage](#)
- [Run a live migration](#)

Migrate a virtual machine in the VMM fabric

7/9/2019 • 6 minutes to read

This article describes how to migrate a VM in System Center - Virtual Machine Manager (VMM).

To perform a migration, you can do any of the following:

- **Run the Migrate VM Wizard:** By using this wizard, you can select a destination virtual machine host for the migration, specify the path that stores the virtual machine files, attach the virtual machine to any of the virtual networks that are found on the selected host, and, if a storage area network (SAN) transfer is available, select a network transfer instead.
- **Drag the virtual machine onto a host:** When you drag a virtual machine to a host, VMM uses automatic placement to place the virtual machine on the most suitable volume on the host. The placement is based on available space.
- **Drag the virtual machine onto a host group:** When you drag the virtual machine to a host group, VMM uses automatic placement to place the virtual machine on the most suitable host that is available in the host group, which is based on the virtual machine requirements and the host ratings. The virtual machine is placed on the most suitable volume on the host. The placement is based on available space. During automatic placement, the [host rating process](#) identifies the most suitable volume on each host.

Note the following before you begin the migration:

- If a correctly configured SAN is available, VMM automatically uses SAN to perform transfers. However, if you use the Migrate Virtual Machine Wizard to perform a transfer, you can override the SAN usage and perform a local area network (LAN) transfer.
- If you migrated a virtual machine that is connected to SAN storage, the virtual machine cannot reconnect to the SAN unless the destination host also has access to that SAN. VMM cannot detect if a virtual machine is connected to a SAN or if the destination host is connected to the same SAN, and therefore cannot provide a warning. You must ensure that the new host is configured to enable the virtual machine to reconnect to the SAN before you migrate the virtual machine.
- To migrate VMs between the hosts with different processors, ensure that you make this exception in the VM **Properties > Processor** by selecting **Allow migration to a virtual machine host with a different processor version**. Else, the migration fails.
- If you change the permissions for a virtual machine through the file system, and then migrate the virtual machine, VMM re-creates the access control list (ACL). All changes that were made outside VMM will be lost.
- If you attempt to migrate a virtual machine on a Hyper-V host soon after you have removed a checkpoint from the virtual machine, the migration might fail. If you attempt a migration before Hyper-V has finished deleting the checkpoint, the migration fails, and you must repair the virtual machine by using the **Undo** option. To avoid this issue, you can ensure that the checkpoint has been deleted, or you can wait for Hyper-V to delete it for you. Verify deletion as follows:
 1. In **Virtual Machines** click the virtual machine > **Actions > Stop**.
 2. In Hyper-V Manage, **Status > Merge in progress** indicates that the checkpoint has not been deleted. Wait until this operation has finished before you migrate the virtual machine.

Migrate a virtual machine with the wizard

1. In **Virtual Machines** view, browse to the host, select the VM, and in **Actions** click **Migrate Virtual Machine**.

2. In **Select Host**, select the destination host. You can check the tabs for more details about the host.

- **Details** — Indicates the status of the host, the operating system, and the type and status of virtualization software. Lists the virtual machines on the host.
- **Rating Explanation** — Lists the factors that resulted in a 0 star rating.
- **SAN Explanation or Deployment and Transfer Explanation** — Lists the factors that make a SAN transfer unavailable. In addition, as of System Center Virtual Machine Manager, the Deployment and Transfer Explanation tab provides an explanation if fast file copy cannot be used. Fast file copy is a feature that was introduced in VMM, based on the Windows Offloaded Data Transfers (ODX) feature introduced in Windows Server 2012 R2. For information about ODX, see [Windows Offloaded Data Transfers Overview](#).

NOTE

The Fast file copy feature is not utilized when migrating a VM from Host to Library.

3. In **Select Path** page, accept the default path, or click **Browse** and browse to the folder in which you want to store the configuration files for the virtual machine, and then click **OK**. Note the following:

- If the target host is a part of a failover cluster that has Cluster Shared Volumes (CSV) enabled, you can store the virtual machine on a CSV Logical Units (LUs) and associated Number (LUN) that is already in use by other highly available virtual machines (HAVMs). With CSV, multiple HAVMs can share the same LUN. The migration of one HAVM does not affect others that are sharing the same LUN. VMM also supports multiple HAVMs per LUN for VMware environments that are configured with VMware VMFS LUNs.
- If you selected a path other than a default virtual machine path and want to store other virtual machines on that path, select the **Add this path to the list of host default paths** check box to add the path to the default paths on the host.
- If you use a network transfer, you have the option to specify separate storage locations for each virtual hard disk (.vhd or .vhdx) file for the virtual machine. By default, all .vhd or .vhdx files are stored in the same location that is specified for the virtual machine.
- If SAN transfers are enabled for this deployment, the virtual machine by default is transferred to the host over the SAN. If you do not want to perform a SAN transfer, select **Transfer over the network even if a SAN transfer is available**. If SAN transfers are not available for this deployment, that option is not available.

4. In **Select Networks**, modify the networks, and then attach them to **None** or to any of the virtual networks that are found on the selected host. The networks area lists each of the virtual network adapters that are currently attached to the virtual machine. Network adapters default to **None** (if you selected **None** in the hardware configuration), or to the best matching virtual network according to the network matching rules.

5. In **Select Virtual SAN**, select the applicable virtual SANs from the drop-down list for each listed virtual HBA. Then click **Next**.

6. In **Summary**, review your settings. To start the VM after deployment click **Start the virtual machine immediately after deploying it to the host**. Click **View Script** to view the Windows PowerShell cmdlets that perform the migration.

7. To start migration, click **Move**. Review progress in **Jobs**.

Migrate a VM with drag and drop

1. In **Virtual Machines**, browse to the virtual machine's current host in the navigation pane.

2. Click the VM and, while you hold down the mouse button, drag the virtual machine to either the host of choice or to the host group of choice in the navigation pane.
3. When you release the mouse button, the system attempts to migrate the virtual machine by using one of the following methods:
 - If you dragged the virtual machine to a host, the system evaluates the host's suitability for the virtual machine and attempts to migrate the virtual machine if the host is found suitable.
 - If you dragged the virtual machine to a host group, the system rates each host in the host group and attempts to migrate the virtual machine to the most suitable of those hosts. For the migration to succeed, a virtual machine path must be configured on the host for the recommended volume.

If you encounter difficulties using drag-and-drop, log out of VMM, and then log back in and try again. You can also try restarting the virtual machine, and then try again.

Migrate storage in the VMM fabric

12/7/2018 • 2 minutes to read

This article describes how to migrate storage in the System Center Virtual Machine Manager (VMM) fabric.

Storage migration enables you to move VM files from one storage location to another, on the same VM host. If the virtual machine is running, you can perform a quick storage migration, which results in little or no service outage for users of the virtual machine. If the virtual machine has more than one virtual hard disk, you can specify a separate location for each virtual hard disk (.vhd or .vhdx) file.

[Read this article](#) if you want to run a live migration of VM storage between two locations on a standalone host.

How to migrate storage

1. In **VM's and Services**, click **All Hosts**, and then select the host on which the virtual machine is deployed.
2. Right-click the virtual machine > **Migrate Storage**.

The Migrate Storage Wizard opens at the **Select Path** page. The current location of the VM configuration files is displayed in the **Storage location for VM configuration**, and the current location of each virtual hard disk (.vhd) is displayed in **Disks**.

3. On the **Select Path** page, do the following:
 - In **Storage location for VM configuration** select an existing default virtual machine path on the list. Browse to a location on the host. VMM automatically changes the paths for all virtual disks to the same path that you specified for the configuration files.
 - Type a path. When you type a new path for the configuration files of the virtual machine, VMM does not automatically change the paths for the virtual disks until you click outside of the **Storage location for VM configuration** box.
 - Select the **Add this path to the list of default storage locations on the host** check box, if you selected a path other than an existing virtual machine path, and you want to add the path to the default paths on the host.
 - Specify the configuration file placement options, as follows:
 - a. Select **Automatically place all VHDs with the configuration** to move all of the virtual machine files to the same location.
 - b. Select **Allow VHDs to be placed individually** to move one or more of the virtual machine files to a different location than the location of the configuration files. If you select this setting, in the **Disks** area, type the new path in the **Location** box for each virtual hard disk, or click **Browse** to browse to the location where you want to store the file. Note that if the virtual machine is running, and you change the path for any of the virtual hard drives, you also must specify a new path for the configuration files of the virtual machine, or the migration operation fails. You must enter the new path even if you want to leave the configuration files in their current location. In that case, you can create a new subfolder within the current location of the configuration files, and then select that new location in the **Storage location for VM configuration** box.
4. On the **Summary** page, click **Move** to begin moving the virtual machine files. Review progress in **Jobs**.

Run a live migration in the VMM fabric

3/14/2019 • 5 minutes to read

This article describes how to run a live migration of virtual machines (VMs) or VM storage, in the System Center - Virtual Machine Manager (VMM) fabric. VMM provides live migration support between standalone Hyper-V hosts, or between cluster hosts that have live migration enabled. [Learn more.](#)

Migrate a VM between two standalone hosts

Migrate a virtual machine from one stand-alone Hyper-V host to another stand-alone Hyper-V host. The VM configuration files and virtual hard disk must be located on an SMB 3.0 file share.

1. In **VMs and Services > All Hosts**, click the standalone source host from which you want to migrate.
2. Click the host and in **VMs**, click the running VM that you want to migrate. Start the machine if it's not running.
3. In **Virtual Machine**, click **Migrate Virtual Machine** to start the Migrate Virtual Machine Wizard.
4. In **Select Host**, review the destination hosts and their associated transfer types. The **Live** transfer type appears if both hosts are configured to connect to the same SMB 3.0 file share.
5. Click the destination host where the transfer type is **Live**, and then click **Next**.
6. In **Summary**, click **Move**. To track the job status, open the **Jobs** workspace.
7. To verify that the virtual machine was migrated, check the **VMs** list on the destination host to make sure the VM is running.

Migrate a VM between clusters

Migrate a VM between clusters. Note that the VM temporarily loses its high availability status. Therefore, a host failure during the migration causes the virtual machine to become unavailable. For live migration between clusters, you should use SMB 3.0 file shares as the storage location. Because the storage does not have to be migrated, the time in which high availability status cannot be guaranteed is very short.

1. In **VMs and Services > All Hosts**, click the cluster node from which you want to migrate.
2. In **VMs**, click the running VM that you want to migrate. Start the machine if it's not running.
3. In **Virtual Machine**, click **Migrate Virtual Machine** to start the Migrate Virtual Machine Wizard.
4. In **Select Host**, review the destination hosts and their associated transfer types. The **Live** transfer type is available for any destination cluster nodes that are configured to connect to the same SMB 3.0 file share on which the VM was originally created.
5. Click a node on a different cluster, and then click **Next**.
6. In **Summary**, click **Move**. To track the job status, open the **Jobs** workspace.
7. To verify that the virtual machine was migrated, check the **VMs** list on the destination node to make sure the VM is running.

Migrate storage between two locations on a standalone host

NOTE

You cannot live migrate storage for a shared VHDX file. You can move the other VM files and perform normal live migration. To move the shared VHDX file to another location, you must shut down the VMs and then move the file.

You can run a live migration of VM storage between locations on standalone hosts. You can move the entire virtual

machine, which includes virtual hard disks (VHDs) and configuration information, or move only specific VHDs to a different location.

1. In **VMs and Services > All Hosts**, click the standalone host where the VM is located.
2. In **VMs**, click the running VM for which you want to migrate storage. Start the machine if it's not running.
3. In **Virtual Machine**, click **Migrate Storage** to start the Migrate Virtual Machine Wizard.
4. In **Select Path > Storage location**, click one of the default storage locations on the host. Or click **Browse** to view all possible storage destinations. Click the destination SMB 3.0 file share or location on the local hard disk, and then click **OK**.

If you specify an SMB 3.0 file share in the **Storage location** list, ensure that you use the fully qualified domain name (FQDN) of the destination server in the share path. For example, instead of `\\fileserver1\smbshare`, use `\\fileserver1.contoso.com\smbshare`.

5. Optionally, select **Add this path to the list of default storage locations on the host**, and then click **Next**.
6. On the **Summary** page, click **Move**. Track progress in **Jobs**.

Run live migrations concurrently

You can run live migration on multiple VMs so that two migrations occur at the same time on the same host. Note that:

- You can't select multiple VMs for a live migration. You need to manually start each migration.
- You can specify how many concurrent migrations to run. The default setting is two, which is the number of simultaneous live migration and storage migrations that are enabled in Hyper-V. For example, a host can participate in one outgoing live migration plus one incoming, two outgoing live migrations, or two incoming live migrations.
- Live migrations and live storage migrations are independent. You can perform two live migrations and two live storage migrations simultaneously. VMM considers live virtual machine and storage migration (live VSM) as one live migration and one storage migration.
- You can view concurrent migrations in progress in Hyper-V Manager > **Actions > Hyper-V Settings> Server > Live Migrations** and **Storage Migrations**. In **Jobs**, verify that the migrations occur simultaneously.

Improve live migration speed

On Hyper-V hosts, you can increase live migration speed using compression, by using SMB as the transport, or both. The compression method uses algorithms that reduce the data that is transmitted over the wire. The SMB method can allow for faster data transfer.

By default, faster live migration is enabled to use the compression method. You can disable, enable, or change the method of faster live migration by changing the live storage migration settings, either at the Hyper-V host level, or for each live migration instance.

Change live migration settings as follows:

1. In Hyper-V Manager click **Actions > Hyper-V Settings > Server > Live Migration**, and then click **Advanced Features**.
2. In **Migration settings > Live migration settings**, do one of the following:
 - To disable faster live migration, click **Standard live migration**.
 - To use compression for faster live migration, click **Use compression**.
 - To use SMB for faster live migration, click **Use SMB as Transport**.

Manage roles and permissions in VMM

12/7/2018 • 2 minutes to read

System Center - Virtual Machine Manager (VMM) allows you to manage roles and permissions. VMM provides:

- **Role-based security:** Roles specify what users can do in the VMM environment. Roles consist of a profile that defines a set of available operations for the role, scope which define the set of objects on which the role can operate, and a membership list that defines the Active Directory user accounts and security groups that are assigned to the role.
- **Run As accounts:** Run As accounts act as containers for stored credentials that you use to run VMM tasks and processes.

Role based security

The following table summarizes VMM user roles.

VMM USER ROLE	PERMISSIONS	DETAILS
Administrator role	Members of this role can perform all administrative actions on all objects that VMM manages.	Only administrators can add a WSUS server to VMM to enable updates of the VMM fabric through VMM.
Fabric Administrator (Delegated Administrator)	Members of this role can perform all administrative tasks within their assigned host groups, clouds, and library servers.	Delegated Administrators cannot modify VMM settings, add or remove members of the Administrators user role, or add WSUS servers.
Read-Only Administrator	Members of this role can view properties, status, and job status of objects within their assigned host groups, clouds, and library servers, but they cannot modify the objects.	The read-only administrator can also view Run As accounts that administrators or delegated administrators have specified for that read-only administrator user role.
Tenant Administrator	Members of this role can manage self-service users and VM networks.	<p>Tenant administrators can create, deploy, and manage their own virtual machines and services by using the VMM console or a web portal.</p> <p>Tenant administrators can also specify which tasks the self-service users can perform on their virtual machines and services.</p> <p>Tenant administrators can place quotas on computing resources and virtual machines.</p>
Application Administrator (Self-Service User)	Members of this role can create, deploy, and manage their own virtual machines and services.	They can manage VMM using the VMM console.

Run As accounts

There are different types of Run As accounts:

- **Host computer accounts** are used to interact with virtualization servers.
- **BMC accounts** are used to communicate with the BMC on hosts for out-of-band management or power optimization.
- **External account** are used to communication with external apps such as Operations Manager.
- **Network device accounts** are used to connect with network load balancers.
- **Profile accounts** are used in Run As profiles when you're deploying a VMM service, or creating profiles.

Note that:

- VMM uses the Windows Data Protection API (DPAPI) to provide operating system level data protection services during storage and retrieval of the Run As account credentials. DPAPI is a password-based data protection service that uses cryptographic routines (the strong Triple-DES algorithm, with strong keys) to offset the risk posed by password-based data protection. [Learn more](#).
- When you install VMM, you can configure VMM to use Distributed Key Management to store encryption keys in Active Directory.
- You can set up Run As accounts before you start managing VMM, or you can set up Run As accounts if you need them for specific actions.

Next steps

- [Set up user roles](#)
- [Set up run as accounts](#)

Set up user roles in VMM

6/6/2019 • 3 minutes to read

This article describes how to set up System Center - Virtual Machine Manager (VMM) user roles.

Before you start

- [Learn more](#) about user roles.
- Ensure you have the right permissions to create the role, or to add users to it.
 - **Administrator role:** Administrators can add and remove users
 - **Delegate Administrator role:** Administrators can create the role. Delegated Administrators can create Delegated Administrator roles that include a subset of their scope, library servers, and Run As accounts.
 - **Read-only Administrator role:** Administrators can create the role. Delegated Administrators can create Read-only Administrator roles that include a subset of their scope, library servers, and Run As accounts.
 - **Virtual Machine Administrator role**(applicable for VMM 2019 and later): Administrators can create the role. Delegated Administrator can create VM administrator role that includes entire scope or a subset of their scope, library servers and Run-As accounts.
 - **Tenant Administrator role:** Administrators and Delegated Administrators can create this role.
- The Administrator role is created by default when you install VMM. The user who performs the installation, and all domain users in the local Administrators group on the server, are added to the Administrator role. You can add or remove members in the role properties.

Create a role

1. Click **Settings > Create > Create User Role**.
2. In the **Create User Role Wizard**, enter a name and optional description for the role, and then click **Next**.
3. In **Profile** page, select the role, and then click **Next**.
4. In **Members**, click **Add** to add user accounts and Active Directory groups to the user role. Add the members in **Select Users, Computers, or Groups**, and then click **Next**.
5. In **Scope**, select the private clouds or host groups that the members of the role can use. Then click **Next**.
6. If one or more **Quotas** pages appear (based on whether you selected private clouds on the previous wizard page), review and specify quotas as needed for each private cloud. Otherwise, skip to the next step. Read-only Administrators can only view items in this defined scope.

To set quotas for the combined use of all members of this user role, use the upper list. To set quotas for each individual member of this user role, use the lower list. By default, quotas are unlimited. To create a limit, clear the appropriate check box under **Use Maximum** and then, under **Assigned Quota**, select a limit. When you have completed all settings, click **Next**.

7. If the **Library servers** page appears, add one or more library servers.
8. In **Networking** click **Add** to add the VM networks that the members of this role can use. Then click **Next**.
9. In **Resources**, click **Add** to add resources. In **Specify user role data path**, click **Browse** to specify a library path that members of this user role can use to upload data. Then click **Next**.

10. In **Permissions** page, select global actions, and any cloud-specific actions that you want to allow members of this role to perform. Click **Next**.
11. If the **Run As accounts** page appears, add Run As accounts that you want the members of this role to be able to use. Otherwise, skip to the next step.
12. If the **Quotas for VM networks** page appears, review and specify quotas to limit the number of VM networks that members of this user role can create. Otherwise, skip to the next step.

To limit the combined number of VM networks that can be created by all members of this user role, use the upper setting. To limit the number of VM networks that can be created by each individual member of this user role, use the lower setting.

13. In **Summary** page, review the settings, and click **Finish** to create the role. Verify the role appears in **Settings > Security > User Roles**.

Create Run As accounts in VMM

12/7/2018 • 2 minutes to read

This article describes how to create and manage Run As accounts in System Center - Virtual Machine Manager (VMM).

A Run As account is a container for a set of stored credentials. In VMM a Run As account can be provided for any process that requires credentials. Administrators and Delegated Administrators can create Run As accounts. [Learn more](#) about different types of Run As accounts.

1. Click **Settings**, and in **Create** click **Create Run As Account**.
2. In **Create Run As Account** specify name and optional description to identify the credentials in VMM.
3. In **User name** and **Password** specify the credentials. The credentials can be a valid Active Directory user or group account, or local credentials.
4. Clear **Validate domain credentials** if you don't need it, and click **OK** to create the Run As account.

Set up self-service in VMM

12/7/2018 • 9 minutes to read

This article describes how to set up self-service in System Center - Virtual Machine Manager (VMM).

VMM offers a number of options for self-service users:

- **Virtual machines/Services:** Users can deploy their virtual machines and services to private clouds. A private cloud can be assigned to multiple self-service user roles. Role-level quotas for each self-service user role with the private cloud in scope are used to allocate cloud compute and storage capacity. Member-level quotas set individual limits for members of the self-service user role.
- **Virtual hard disks:** Users can deploy VMs from VHDs as well as templates.
- **Templates/Profiles:** Users can create their own templates and profiles. The **Author** action for a self-service user role providing these authoring rights to create hardware profiles, guest operating system profiles, app profiles, SQL Server profiles, VM templates, and service templates. Note that these resources can be created by a user with the self-service role, and shared with other members of the self-service user role.

Self-service users use the VMM console (or PowerShell) to create and manage VMs, services etc. In the VMM console self-service users can view status, resource usage, jobs, and PRO tips (if enabled) for their own VMs and services. They can view available capacity and quota usage in their private clouds. They can't see host groups, hosts, library servers and shares, or network and storage configuration settings.

You set up self-service in VMM as follows:

1. Create a self-service user role. Specify actions that the role can perform, assign resources to the role, and configure Run As accounts that self-service users can use when interacting with VMM.
2. Set up the VMM library. Assign a library share on which resources available to self-service users will reside. In addition, set up a share so that self-service users can share their resources with other users.

Set up a self-service user role

1. Click **Settings > Create > Create User Role**.
2. In the **Create User Role Wizard**, enter a name and optional description for the role, and then click **Next**.
3. In **Profile** page, select **Self-Service User**, and then click **Next**.
4. In **Members**, click **Add** to add user accounts and Active Directory groups to the role. Then click **Next**.
5. In **Scope**, select at least one private cloud that members of the role will use. Then click **Next**.
6. In **Quotas**, set quota for each private cloud. Each quota sets an individual limit for each member of the user role. If you want all role members to share overall quotas, create a security group in Active Directory, and assign that group to the user role. Supported quota types include:
 - **Virtual CPUs:** Limits the total number of VM CPUs that can be consumed from the private cloud.
 - **Memory (MB):** Limits the amount of VM memory that can be consumed from the private cloud.
 - **Storage (GB):** Limits the amount of VM storage that can be consumed from the private cloud.
 - **Quota (points):** Sets a quota on VMs deployed on the private cloud based on total quota points assigned to the VMs via their VM templates.
 - **Virtual machines:** Limits the total number of VMs that can be deployed on a private cloud. include
7. In **Resources**, click **Add** to add resources that the role can use. You can assign hardware profiles, OS

profiles, VM templates, app profiles, SQL Server profiles, and service templates that can be used when creating VMs and services.

8. In **Specify user role data path**, click **Browse** to specify a library path that members of this user role can use to upload and share their own data. Then click **Next**.
9. In **Actions**, select the actions that users are allowed to perform.
 - **Author**: Users can author templates and profiles, including hardware profiles, operating system profiles, application profiles, SQL Server profiles, virtual machine templates and service templates.
 - **Checkpoint**: Users can create, edit, and delete checkpoints for their own VMs, and to restore a VM to a previous checkpoint. VMM doesn't support checkpoint actions on services.
 - **Checkpoint (Restore only)**: Users can restore their own VMs to a checkpoint but can't create, edit, and delete checkpoints.
 - **Deploy**: Users can deploy virtual machines and services from templates and virtual hard disks that are assigned to their role. They can't author templates and profiles.
 - **Deploy (from template only)**: Users can deploy VMs and services from templates only. They don't have authoring rights.
 - **Local Administrator**: Users can be Local Admins on their own VMs. You must enable **Local Administrator** on any User role that has the **Deploy (From template)** enabled, so that those users can set the Local Admin password during VM and service deployment. Users will the Deploy action don't need this to set credentials.
 - **Pause and resume**: Users can pause and resume their own VMs and services.
 - **Receive**: Users can use resources that are shared by members of other self-service user roles.
 - **Remote connection**: Users can connect to their VMs from the VMM console or App Controller.
 - **Remove/Save**: Users can remove or save their VMs.
 - **Share**: Users can share resources that they own with other self-service user roles. Sharable resources include hardware profiles, operating system profiles, application profiles, SQL Server profiles, virtual machine templates, virtual machines, service templates, and services. A self-service user must be the owner of a resource to share it. For a user role to use the resources, it must have the **Receive** action.
 - **Start/Stop**: Users can start and stop their own VMs and services. **Store and redeploy**: Users can store their own virtual machines in the VMM library, and redeploy those virtual machines. Virtual machines stored in the library do not count against a user's virtual machine quota. VMM doesn't support storing services.
10. If the **Run As accounts** page appears, add Run As accounts that you want the members of this role to be able to use in the actions to create VMs and services. Then click **Next**.
11. In **Summary** page, review the settings, and click **Finish** to create the role. Verify the role appears in **Settings > Security > User Roles**.

After you create the role, you can modify its settings on the properties page.

Prepare the VMM library for self-service

Self-service users with the required permissions can access the VMM library. Users with the Author action can create templates and profiles in the library. They can also share those templates and profiles with other self-service users. In order for self-service users to interact with the library you need to prepare the following:

- **Read-only library shares**: To share physical resources such as VHDs and ISO images with self-service users, you set up read-only library shares for private clouds, and add the resources to the path. The resources are then available for self-service users that have the private cloud in their scope. You could also store resources such as Application Frameworks on these shares to enable self-service users to configure templates and profiles with scripts.

- **Self-service user data paths:** Set up user data paths on self-service roles to provides a place where members of the role can upload and share their own resources. For example a path might store app packages for services deployed by a self-service user role. Read and write permissions for the path are controlled through the file system. VMM discovers all paths that the current self-service user can access. These data paths must be on a library share.

Before you start

All of these procedures must be performed by a VMM administrator. Delegated administrators can add library shares on library servers that are in the scope of their user role, can configure read-only library shares on private clouds that they created, and can configure user data paths on self-service user roles that they created. Only members of the local Administrators group can grant access permissions on their user data paths.

Create read-only library shares

1. Create a shared folder to store resources. The folder will include read-only library shares for private clouds, and user data paths for self-service user roles. We recommend that you crate the folder near your default library share so that it's easy to access when you're managing the library. For example C:\ApplicationData\Virtual Machine Manager Cloud Resources.
2. In the shared folder, create a folder to store the \ApplicationFrameworks resources in case you want to use them. For example C:\ApplicationData\Virtual Machine Manager Cloud Resources\ApplicationFrameworks. Share the folder so that you can add it as a library share. Note that the shared folder can't be in the default library share path. You can't add a library shares that's in the path of an existing library share.
3. Copy the \ApplicationFrameworks folder from the default library share to the share you created for private cloud resources.
4. Add the share to the VMM library. In **Library > Library Server > Add Library Share**, select each shared folder you want to add to the library. Verify that the share is added in **Library Servers**.
5. To add the read-only share to a private clouds open VMs and Services > clouds, and select the private cloud you want to update.
6. In the cloud, click **Folder > Properties > Library > Read-only library shares > Add**.

Enable self-service users to share resources

To enable self-service users with the Author action to share resources they create you need to create a folder to store shared resources, and then enable resource sharing for the self-service user role.

Create a folder to share user resources

Configure a user data path for the self-service user role, and grant read/write permission on the folder.

1. Create a folder to store all resources that will be shared by self-service users. For example, C:\ProgramData\Virtual Machine Manager Cloud Resources\Self-Service User Data.
2. Within that folder, create a subfolder to store resources for the self-service user role. For example: C:\ProgramData\Virtual Machine Manager Cloud Resources\Self-Service User Data\Finance Service Managers.
3. Then within that subfolder, create a third-level subfolder to store all the application packages for all releases of the virtual application that you will use in this scenario. For example: C:\ProgramData\Virtual Machine Manager Cloud Resources\Self-Service User Data\Finance Service Managers<MyApplication>.
4. In that subfolder, create a fourth-level subfolder to store the application package for the first release of the service. For example: C:\ProgramData\Virtual Machine Manager Cloud Resources\Self-Service User Data\Finance Service Managers<MyApplication>\MyApplication v1 >.

Each time you update and re-sequence an application by using Server App-V, you will need to store the new application package in a separate folder.

5. To enable members of the self-service user role to access the resources and upload their own resources to the folder, grant all members read/write permission on the folder.
6. If needed, share the folder that contains user data for all self-service user roles, and then add the share to the VMM library. To be assigned to a self-service user role, a user data path must be on a library share.
7. Configure the path for a self-service user role as follows:
 - a. In **Settings > Security > User Roles**, click the self-service user role.
 - b. In the **User Role** group, click **Properties > Resource**.
 - c. Browse and select the folder that will hold the shared resources. After you save the changes, the data path is added to the library. Verify the path in **Library > Self-Service User Content**.

Enable sharing for self-service users

To share a resource with a member of another self-service user role, you need the following:

- The self-service user who shares the resource must be the owner of the resource.
- The resource owner must belong to a self-service user role that has been assigned the Share action.
- The resource receiver must belong to a self-service user role that has been assigned the Receive action.

Enable resource sharing as follows:

1. Click **Settings > Security > User Roles**, and click the self-service user role for which you want to enable resource sharing.
2. In the **User Role** group, click **Properties**.
3. In **Actions**, select **Share**, and then click **OK**. Members of this self-service user role can now share their own resources with members of any self-service user role that has the **Receive** action assigned to it.
4. To configure a user role with the **Receive** action, select the role > **Properties > Action**, and select **Receive**.

Work with VMM as a self-service user

12/7/2018 • 3 minutes to read

This article describes how to work with System Center - Virtual Machine Manager (VMM) as a self-service user.

Self-service users can interact with VMM to deploy virtual machines and services to private clouds. Depending on your permissions, you can deploy VMs from VHDs and templates, and create and share your own templates and profiles. You interact with VMM using the VMM console (or PowerShell).

Create and deploy virtual machines

- [Learn about](#) creating a VM from an existing VHDs
- [Learn about](#) creating a VM from a blank VHDs
- [Learn about](#) cloning an existing VM
- [Learn about](#) deploying VMs from VM templates
- [Learn about](#) deploying VMs running Linux

Create resources in the VMM library

The VMM library is a file share that includes a catalog of resources that are used to deploy virtual machines and services in the VMM fabric. The library stores:

- File-based resources such as virtual hard disks, ISO images, and scripts, driver files and application packages (SQL Server data-tier applications and Web Deploy).
- Non-file-based resources such as virtual machine templates and service templates that are used to create VMs and services.
- Offline virtual machines are stored in the library.

If you have Author permissions you can create templates and profiles in the library. [Learn more](#).

Share library resources

As a self-service user you can share your library resources with other members of your self-service user role when the following conditions are met:

- You must be the resource owner.
- You must belong to self-service user role that has been assigned the Share action.
- To use the shared resource, the user must belong to a self-service user role that has been assigned the Receive action.

Share resources as follows:

1. Click **Library**, right-click the template or physical resource that you want to share > **Properties**.
2. In **Access > Users**, click **Add**.
3. In **Select Users**, in User, share the resource with another member of your role. In **User Role**, share the resource with other self-service user roles.

Import library resources

We recommend you use the method described in this procedure to import and export file-based resources to and from the VMM library. Note the following:

- You can import resource to the user role data path in the Self Service User Content node in the library.

- You can export resources from your user role data path, or private cloud library.
- To import and export your self-service user role needs Author permissions.

Import resources as follows:

1. Click **Library > Import > Import Physical Resource**.
2. Click **Add custom resource** to import a folder with the .CR extension and its contents. Alternatively, you can select a folder without a .CR extension that contains one or more files of a supported file type.
 - If you select a folder without a .CR extension, only the files of a supported file type will appear in the VMM library.
 - However, if you use Windows Explorer to access the library share, you can access all files in the folder depending on the file and share permissions that are configured outside VMM.
 - If the folder with .CR extension contains more than one hundred file to be imported, it is recommended that you zip the files before the import. This will improve performance.
3. Click **Add resource** to import one or more files to another library location.
4. When you are finished, click **Import**. Check that the resources are listed **Self Service User Content > user role data path > Self Service User Objects**.

Export library resources

1. Click **Library > Export > Export Physical Resource**.
2. Click **Add** and select the physical resources you want to export. Then click **OK**.
3. In **Specify a destination for the export files**, click **Browse** and select a destination folder. Then click **OK**.
4. When you're finished, click **Export**.

Next steps

[Provision VMs](#)

Monitor VMM

12/7/2018 • 4 minutes to read

You configure monitoring and reporting in System Center - Virtual Machine Manager (VMM) as follows:

- **VMM jobs:** In the VMM console you can monitor the status of processes and operations with VMM jobs.
- **Monitoring in Operations Manager:** To monitor the health and status of VMM servers and the VMM fabric you integrate VMM with System Center Operations Manager
- **Reporting in Operations Manager:** After VMM is integrated with Operations Manager, you can create and view VMM reports

Monitor with VMM jobs

A VMM job is created for any action that changes the status of a managed object.

- Jobs are composed of steps that are performed sequentially to complete an action. Simple jobs, such as stopping a virtual machine, contain only one step. More complex jobs, or running a wizard, might contain multiple jobs, or job groups.
- Jobs are tracked in the **Jobs** view in the VMM console. The **Details** tab in the **Jobs** view shows the status of each step in a job.
- Each job is represented by one or more VMM PowerShell cmdlets. You can perform almost any VMM task with [PowerShell](#). Each wizard includes a **View Script** button on the **Summary** page, that displays the cmdlets that will perform the job that you just requested. You can save Windows PowerShell scripts to the VMM library, and view and run them in **Library** view.
- Each VMM job is independent, and doesn't depend on the status of another job. For example, if you're running jobs to add multiple host servers, the failure to add one host does not affect the remaining jobs. When a job completes, an audit record is saved that lists the changes that the job made to the VMM object. You can view the audit record in **Jobs > Details > Change Tracking**.
- Jobs are started automatically when you perform tasks in the VMM console, or using PowerShell. You can cancel some running jobs, but others, including adding hosts, and system jobs, can't be cancelled after they've started running.
- You can generally restart failed jobs, but note that:
 - If multiple jobs place a VM into a failed state, only the most recent job can be restarted.
 - For jobs that run for a long time, such as virtual machine migration, intermediate results are stored periodically while the job is running, and the Restart action attempts to resume the job from the last known state. All other jobs start from the beginning.

Monitor with Operations Manager

You can monitor VMM health and status in Operations Manager by installing the VMM management pack which provides a number of dashboards in the Operations Manager console.

DASHBOARD	DETAILS
Virtual Machine Dashboard	Monitors the health of virtual machines It shows information about discovered VMs in the VMM fabric. You can view alerts, properties, host, and performance information for the VM.

DASHBOARD	DETAILS
VMM Host Dashboard	<p>Monitors the health of virtualization hosts discovered in the VMM fabric.</p> <p>It shows information about the host properties, status, VMs, alerts, and performance.</p>
Fabric Health Dashboard	<p>Monitors the health of VMM private clouds</p> <p>For each cloud the dashboard monitors the state of host groups, the compute properties of the cloud (CPU, network adapters etc), storage information such as pools, file share, LUNs, and network monitoring.</p> <p>A fabric monitoring dashboard diagram view shows the status for each piece of the fabric.</p> <p>The dashboard can be scoped to physical or virtual resources, or to a particular cloud tenant.</p>

[Learn more](#) about integrating VMM with Operations Manager.

Report with Operations Manager

After you've connected VMM to Operations Manager you can view and create reports. VMM provides these default reports:

REPORT	DETAILS
Capacity Utilization	<p>Details usage for VM hosts and other objects.</p> <p>Provides an overview of how capacity is being used in your datacenter, and helps you to make resources decisions for supporting your VMs.</p>
Host Group Forecasting	<p>Predicts host activity based on history of disk space, memory, disk IO, network IO, and CPU usage.</p> <p>To use the forecasting reports, SQL Server Analysis Services must be installed on the Operations Manager Reporting server.</p>
Host Utilization	<p>Shows the number of virtual machines that are running on each host and average usage, along with total or maximum values for host processors, memory, and disk space.</p>
Host Utilization Growth	<p>Shows the percentage change in resource usage and the number of virtual machines that are running on selected hosts during a specified time period.</p>
Power Savings	<p>Shows how much power is saved through power optimization.</p> <p>You can view total hours of processor power saved for a date range and host group, as well as detailed information for each host in a host group.</p>
SAN Usage Forecasting	<p>Predicts SAN usage based on history.</p>

REPORT	DETAILS
Virtual Machine Allocation	Provides information about allocation of virtual machines.
Virtual Machine Utilization	Provides information about resource utilization by virtual machines, including average usage and total or maximum values for virtual machine processors, memory, and disk space.
Virtualization Candidates	<p>Helps identify physical computers that are good candidates for conversion to virtual machines.</p> <p>You can use this report to identify little-used servers and display average values for a set of commonly requested performance counters for CPU, memory, and disk usage, along with hardware configurations, including processor speed, number of processors, and total RAM.</p> <p>You can limit the report to computers that meet specified CPU and RAM requirements, and you can sort the results by selected columns in the report.</p>

View reports

You can view reports in the Reporting workspace in System Center Operations Manager, or by using a web browser and entering this address: `http://ReportingServerName:port/reports`. You can optionally specify "https" instead of http.

- `ReportingServerName` is the name of the Operations Manager reporting server.
- `port` is 80 for HTTP and 443 for HTTPS
- `reports` is the reporting server virtual directory, by default: **reports**

Integrate VMM with Operations Manager for monitoring and reporting

5/6/2019 • 6 minutes to read

This article describes how to integrate System Center - Virtual Machine Manager (VMM) with System Center Operations Manager to monitor the health of virtualization hosts, virtual machines, and other resources in the VMM compute fabric.

- Operations Manager monitors the VMM server and all hosts and virtual machines in the VMM fabric.
- An Operations Manager management group can monitor multiple VMM instances. The VMM fabric can only be monitored by a single Operations Manager management group.
- The Fabric Health Dashboard is imported to Operations Manager with the VMM management packs. It shows a detailed overview of the health of the VMM fabric and private clouds. You can link to other dashboards to drill into network and storage monitoring.

Deployment summary

You set up Operations Manager with VMM as follows:

1. Verify prerequisites.
2. Install the Operations Manager console on the VMM server so you can monitor VMM from the server.
3. Install Operations Manager agents on the VMM management server and all hosts under management by VMM.
4. Download the latest management pack.
5. Run the integration wizard to integrate VMM and Operations Manager. The wizard does the following:
 - Imports the VMM management packs into Operations Manager.
 - Optionally enables Performance and Resource Optimization (PRO). PRO information is provided by Operations Manager, and can be VMM to optimize performance. You can map specific Operations Manager alerts to remedial actions in VMM. For example, you could migrate VMs to a different host after a hardware issue. In addition, with PRO enabled, Operations Managers can detect resource issues or hardware failures in the virtualization infrastructure.
 - Optionally enables maintenance mode. VMM can place hosts in maintenance mode for servicing. When a host is in maintenance mode, VMM uses live migration to move VMs to another location, and doesn't place new VMs on the host. If maintenance mode is enabled for Operations Manager monitoring, when a host is placed into maintenance mode in VMM, Operations Manager also places it in the same mode. In maintenance mode, the Operations Manager agent suppresses alerts, notifications, and state changes, so that the host isn't monitored while regular hardware and software maintenance activities are in progress.
 - Enables support for SQL Server Analysis Services (SSAS) and the reporting capabilities provided by SSAS.

Before you start

- Make sure you're using a supported version of Operations Manager (running on Systems Center 2016 or a later version).

- Operations Manager must use SQL Server 2012 SP2, SQL Server 2014, or SQL Server 2016 with reporting services enabled. To use the forecasting reports, SQL Server Analysis Services must be installed on the Operations Manager reporting server. The SSAS instance name should match the SQL Server Reporting Services (MSSQLSERVER).
- The version of the Operations Manager operations console that is installed on the VMM management server must match the version of Operations Manager with which you intend to integrate. The Operations Manager agent version agent should be supported by the Operations Manager version.
- Ensure that the version of Windows PowerShell that's on all Operations Manager management servers is the most recent version supported by that version of Operations Manager. To determine which version of Windows PowerShell is on a server, run **Get-Host | Select-Object Version**
- Make sure that port 5724 is open between the VMM and Operations Manager servers.
- You need VMM admin permissions to run the integration wizard.
- VMM monitoring packs support up to 400 hosts, with up to 8000 virtual machines.

Install the Operations Manager console

1. Run Operations Manager Setup on each VMM management server.
2. In the wizard, select **Operations** Console to install it.

Install the Operations Manager agent

1. Install the Operations Manager agent on each VMM server, and on each host managed in the VMM fabric.
2. You can install the agent from the Operations Manager Operations console, from Operations Manager setup, or from the command line. To understand the different methods for agent deployment, review [Managing discovery and agents](#) section in Operations manager documentation.

Download the management packs

1. [Download](#) the latest version of the VMM management packs from the Microsoft download center.
2. Open the management packs folder on the VMM server. The default location is C:\Program Files\Microsoft System Center \Virtual Machine Manager\ManagementPacks. Back up the existing .mp files.
3. Extract the new .mp files to the management pack folder, overwriting the existing .mp files.

Run the integration wizard

Run the wizard to connect the VMM server to the Operations Manager server, and import the VMM management pack to Operations Manager.

1. In the VMM console, click **Settings > System Center Settings > Operations Manager Server > Properties**.

NOTE

If the Operations Manager connection has already been established, clicking **Properties** opens the **Operation Manager Settings** dialog box. If this dialog box appears and does not describe the correct connection, remove the current connection before you enter the correct information.

2. In **Introduction**, click **Next**.
3. In **Connection to Operations Manager**, specify the Operations Manager server name, and select an account to use to connect to it. You can use the VMM server service account or specify a Run As account. This account must be a member of the Operations Manager Administrator role.

4. Select **Enable Performance and Resource Optimization (PRO)** if required.
5. Select **Enable maintenance mode integration with Operations Manager**, if desired. Then click **Next**.

When hosts are placed in maintenance mode using the VMM management server, Operations Manager places them in maintenance mode as well. In this mode, the Operations Manager agent suppresses alerts, notifications, rules, monitors, automatic responses, state changes, and new alerts.

6. Enter credentials for Operations Manager to connect with the VMM management server, and then click **Next**. This account will be added to the Administrator user role in VMM.
7. Review the information in the **Summary** page, and click **Finish**. You can view the status of the new connection in the **Jobs** workspace.
8. With **System Center Settings** still selected, in the results pane, right-click **Operations Manager Server**, and then click **Properties**. In **Operation Manager Settings > Details > Connection Status**, confirm that the connection is **OK**.

If you later remove a connection to an Operations Manager server, this doesn't remove the VMM management packs from the server, but the connector is removed.

Monitor VMM in the Operations Manager console

You can monitor VMM processes and state in any of the VMM dashboards that appear in the Operations Manager console - the fabric dashboard, VM dashboard, or VMM host dashboard.

1. In the Operations console for Operations Manager, click **Monitoring**.
2. To monitor VMs, click **Virtual Machine Manager**. This dashboard includes health and performance information for virtual machines, hosts, and VMM servers.
3. To monitor VMM hosts, click **VMM Host Dashboard**. You can monitor the health of virtualization hosts discovered in the VMM fabric, and get information about the host properties, status, VMs, alerts, and performance.
4. To monitor a private cloud click **Fabric Health Dashboard**. For each cloud the dashboard monitors the state of host groups, the compute properties of the cloud (CPU, network adapters etc), storage information such as pools, file share, LUNs, and network monitoring. You can scope the dashboard to physical or virtual resources, or to a particular cloud tenant.
5. In **Virtual Machine Manager Views** you can view graphic representations of managed systems. After you connect the VMM and Operation Manager servers for the first time it can take several hours until the diagrams are available in the console.

Next steps

- [Run VMM reports](#) in Operations Manager

Add an Azure subscription in VMM

5/16/2019 • 5 minutes to read

You can add Microsoft Azure subscriptions to System Center - Virtual Machine Manager (VMM), and perform basic actions on Azure instances in the subscriptions.

- For each Azure subscription you add, you can use a console to see all role instances in all Deployment Groups in that subscription.
- If you are already managing your on-premises virtual machines in VMM, you can use this feature to perform some very basic actions on Azure instances without leaving the VMM console. You can:
 - Add and remove Azure subscriptions in the VMM console.
 - Get a list view of information about role instances in deployment subscriptions. Instances can be manually refreshed.
 - Perform basic actions on the instances, including start, stop, shutdown, and restart.
 - Connect via RDP
- This feature isn't designed to provide feature parity with the Azure portal. It provides a small subset of the features to simplify management. You can't manage the Azure subscription, deploy instances, manage Azure storage and networks, migrate on-premises VMs to Azure, or view the dashboard and performance monitoring summaries.
- Certificates and subscription setting information is stored in the registry under HKEY_CURRENT_USER. The information is log in-specific, and visible on a per-machine, per-login basis. If you log into VMM with a shared account (not recommended), any subscription added under that shared account will be exposed to all users.
- Because information is log in-specific, Azure subscriptions added by one admin user aren't automatically visible to all other VMM admins. You need to set up the subscription for each VMM admin that needs to access it.

Before you start

Here's what you need to add an Azure subscription in VMM:

REQUIREMENT	DETAILS
Azure subscription	You need at least one Azure subscription to add it to the VMM console.
Internet connectivity	The computer on which you install the feature must be able to connect to the Azure subscription.
Service administrator	You need to be at least a service administrator for the subscription. You need this for access to the management certificate information that's required.

REQUIREMENT	DETAILS
Management certificate	<p>The subscription must have a management certificate associated with it if you are managing Classic VMs only. So that VMM can use the service management API in Azure. Learn more about service certificates. Make note of the subscription ID and the certificate thumbprint.</p> <p>Certificates must be x509 v3 compliant.</p> <p>The management certificate must be located in the local certificate store on the computer on which you add the Azure subscription feature.</p> <p>The certificate should also be located in the Current User \ Personal store of the computer running the VMM console.</p>

NOTE

To enable management of both Classic and ARM-based VMs, the subscription must have **Active Directory-based** authentication associated with it. Create an Azure AD application using Azure portal and make note of Directory ID, Application ID and Key.

NOTE

To enable management of both Classic and ARM-based VMs, the subscription must have **Active Directory-based** authentication associated with it. Create an Azure AD application using Azure portal and make note of Directory ID, Application ID and Key.

Set up the Add Azure Subscription feature

1. Make sure you've complied with requirements in the previous section, and ensure that the option to add a subscription appears in the VMM console > **VMs and Services** > **Azure** > **Add Subscription**.
2. If you want to check for a certificate, log into the Azure portal with your account, and click **Settings** > **Management Certificates**. Note the certificate thumbprint if you've haven't already.
3. In the VMM console, click **VMs and Services** > **Azure** > **Add Subscription**, and fill in the subscription details. In the Subscriptions node you can add and remove subscriptions, and refresh VM lists. You can also start, stop, shut down, and restart a VM, and connect to it over RDP.

Manage Azure VMs

If you are already managing your on-premises virtual machines in VMM, you can use this feature to perform some very basic actions on Azure instances, without leaving the VMM console. You can:

- Add and remove one or more Azure subscriptions by using the VMM console.
- See a list view with details and status of all role instances in all deployments in that subscription.
- Manually refresh the list of instances.
- Perform the following basic actions on the instances:
 - Start
 - Stop
 - Shutdown
 - Restart
 - Connect via RDP

What you can't do with this feature

This feature is not intended to provide feature parity with the Microsoft Azure Management Portal. The functionality of this feature is a minor subset of the features at <https://manage.windowsazure.com>, but you can view your instances and do other basic actions, to simplify everyday tasks and help make management easier.

You cannot:

- Manage your Azure subscription
- Deploy instances to Azure
- Migrate on-premise virtual machines to Azure
- Manage Azure Storage
- Manage Azure Networks
- See the Dashboard Summary view
- See the Performance Monitoring Summary

You can add Microsoft Azure subscriptions to System Center 2019 - Virtual Machine Manager (VMM), by creating an Azure profile.

With VMM 2019, using Azure profile, you can define the intended usage of the profile. Currently, Azure-VMM integration scenario supports the following:

- **Azure VM Management** – perform basic actions on Azure VM instances without leaving the VMM console.
- **Azure Update Management** – install the update on the VMs managed by VMM.

Before you start

Here's what you need to add an Azure profile for Azure VM management:

REQUIREMENT	DETAILS
Azure subscription	You need at least one Azure subscription to add it to the VMM console.
Internet connectivity	The computer on which you install the feature must be able to connect to the Azure subscription.
AD Authentication	<p>To enable management of both Classic and Azure Resource Manager based VMs, the subscription must have Active Directory-based authentication associated with it.</p> <p>Create an Azure AD application using Azure portal and make a note of the directory ID, application ID and Key.</p> <p>Assign application to Classic VM contributor and VM contributor roles by using <i>Subscription – Access Control (IAM) – Add</i>.</p>

Here's what you need to create an Azure profile for Azure Update Management:

REQUIREMENT	DETAILS
Azure subscription	<p>You need Azure Automation Subscription with Update Management solution enabled.</p> <p>Create Automation Account and Enable Update Management Solution.</p>

REQUIREMENT	DETAILS
Internet connectivity	The computer on which you install the feature must be able to connect to the Azure subscription.

Create Azure Profile

Follow these steps:

1. In VMM console, go to **Library > Create > Azure Profile**

The screenshot shows the 'Create Azure Profile' wizard. The 'General' tab is active. The 'Name' field is 'Test'. The 'Description' field is empty. The 'Cloud Provider' is 'Public Azure'. The 'Subscription ID' is '123ADC-FGD'. The 'Profile Usage' dropdown menu is open, showing three options: 'Azure VM Management', 'Azure VM Management', and 'Azure Update Management'. The 'Next' button is highlighted.

2. Under **Profile usage** drop-down menu, select **Azure VM Management** or **Azure update management**. Based on the selection, the next page seeks authentication information for the subscription ID entered.

NOTE

- You can share Azure profile with Self Service Users (SSUs) by adding them as members in the wizard.
- You can view the list of all Azure profiles from **Library > Profiles > Azure Profiles**. Select an Azure profile from the list to view detailed information of this profile, under **General Information** pane.

Azure update management

4/23/2019 • 2 minutes to read

This article provides information about Azure update management feature in System Center 2019 - Virtual Machine Manager (VMM).

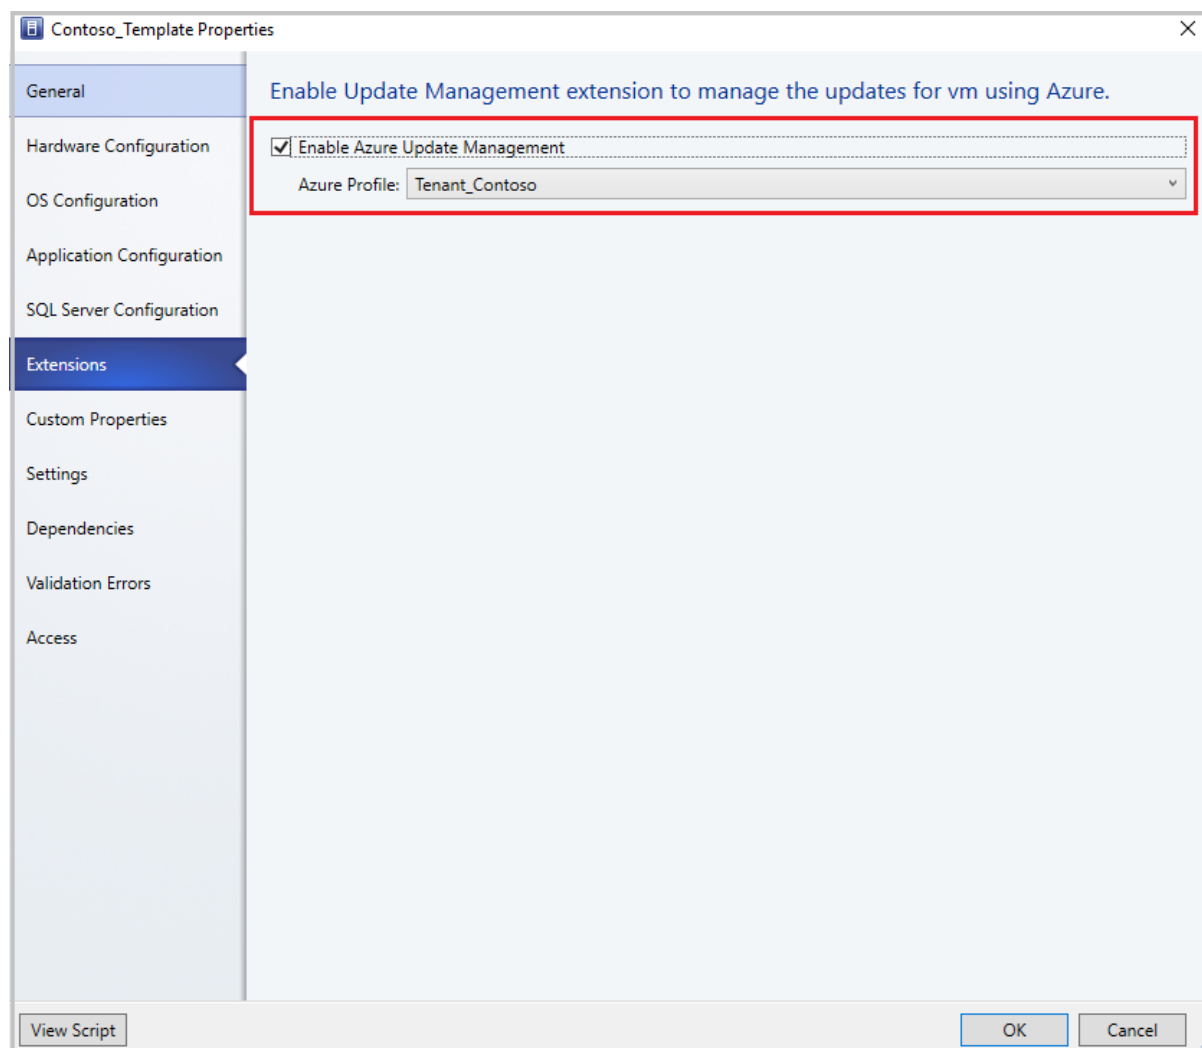
Using Azure update Management feature, you can manage updates for Virtual Machines (VMs) and Workloads running in a VMM.

Currently, VMM supports update management feature for all new VMs with Windows operating system, and are deployed using a VM template with *Azure Update Management Extension* enabled.

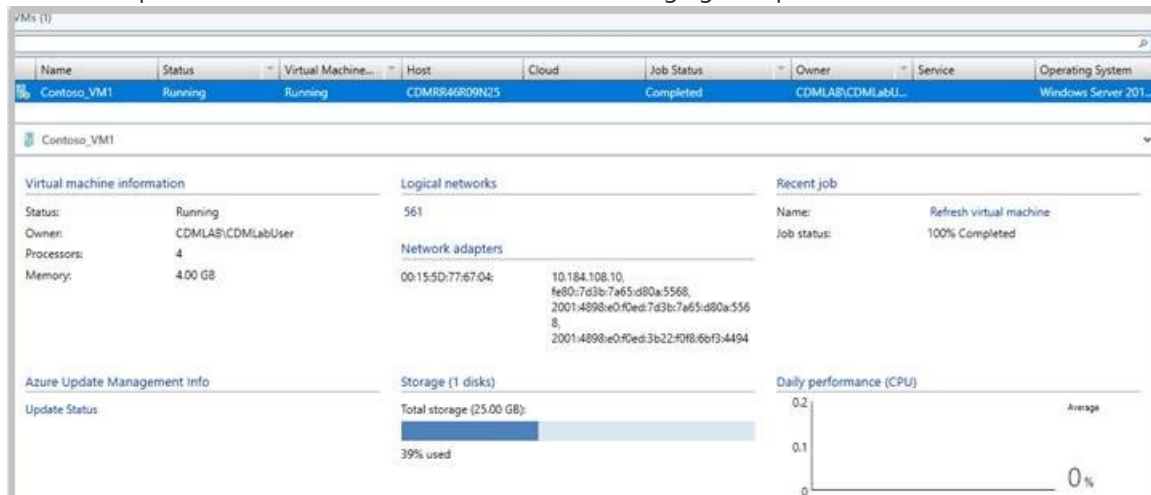
Create a VM template linked to Azure profile

Follow these steps:

1. Create a profile for Azure Update management using steps detailed in [Azure subscriptions article](#)
2. In the **Create VM Template** wizard, select **Source Page** > select **Use an existing VM template or a virtual hard disk stored in the library**.
3. On the **Extension** page, select **Enable Azure Update Management** and select the Azure Profile from the drop-down menu. Click **OK**.



4. Deploy the VMs from the VM template.
5. VMM does the onboarding for the VMs deployed through VM template to the *Azure Update Management* service and provides the link to the Azure console for managing the updates.



6. Click the **Update Status** link under **Azure Update Management info** to assess and deploy the updates for the VM.

NOTE

Azure update management capability also supports service deployments using service templates, the flow is same as the above.

Next steps

[Manage Azure VMs](#)

Manage Azure VMs

3/14/2019 • 2 minutes to read

This article provides information about the *Manage Azure Virtual Machines (VMs)* feature in System Center 2019 - Virtual Machine Manager (VMM).

This feature allows you to perform basic actions on azure instances attached to an azure profile, created for Azure VM management.

NOTE

To perform these actions, you must create [Azure profile](#) with **Profile Usage** selected as **Azure VM Management**. Once an Azure profile for Azure VM Management is created, the VMs deployed on Azure will be accessible on the **VMs and Services** page of VMM console, and will be listed under *Azure Subscriptions*.

You can perform the following actions on Azure instances, without leaving the VMM console.

- Add and remove one or more Azure subscriptions by using the VMM console.
- See a list view with details and status of all role instances in all deployments in that subscription.
- Manually refresh the list of instances.
- Perform the following basic actions on the instances:
 - Start
 - Stop
 - Shutdown
 - Restart
 - Connect via RDP

What is not supported

This feature is not intended to provide feature parity with the Microsoft Azure Management Portal. The functionality of this feature is a minor subset of the features at <https://manage.windowsazure.com> but you can view your Azure instances and do other basic actions, to simplify everyday tasks and management.

With this feature, You cannot:

- Manage your Azure subscription
- Deploy instances to Azure
- Migrate on-premises virtual machines to Azure
- Manage Azure Storage
- Manage Azure Networks
- See the Dashboard Summary view
- See the Performance Monitoring Summary

Manage VMs using Azure AD-based authentication & authorization and region-specific Azure subscriptions

5/30/2019 • 2 minutes to read

This article provides information about how to manage the Azure Resource Manager-based and region-specific Azure subscriptions using System Center - Virtual Machine Manager (VMM).

You can add Microsoft Azure subscriptions to System Center 2016 - Virtual Machine Manager (VMM) and later, and perform required actions. [Learn more](#). The VMM Azure plugin allows management of Azure subscriptions through certificate-based authentication and authorization, and VMs in public Azure region.

VMM 1801 and later supports management of Azure subscriptions through Azure active directory and region-specific Azure subscriptions. (namely, Germany, China, US Government Azure regions).

Management of Azure subscriptions through certificate-based authentication and authorization requires Management certificate. [Learn More](#).

Management of VMs using Azure AD-based authentication and authorization requires Azure AD application.

Before you start

Ensure the following prerequisites:

- **Azure AD application** - to manage VMs by using VMM through AD authentication and authorization, you need to create an Azure AD application, and then provide the following details through VMM Azure plugin:
 - Azure Subscription ID
 - Azure Active Directory ID
 - Azure Active Directory - Application ID & Application Key[Learn more](#) on how to create an Azure AD App.
- **A management certificate** - with the configuration as described in [this article](#).
 - The subscription must have a management certificate associated with it so that VMM can use the service management API in Azure.
 - Make note of the subscription ID and the certificate thumbprint.
 - Certificates must be x509 v3 compliant.
 - The management certificate must be located in the local certificate store on the computer on which you add the Azure subscription feature.
 - The certificate should also be located in the Current User\Personal store of the computer running the VMM console.

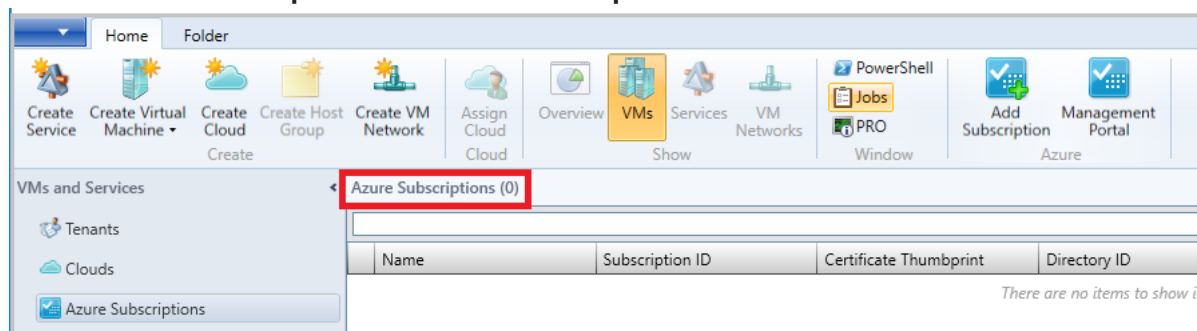
NOTE

The certificate is required only if you choose to use certificate-based authentication to manage your Azure subscription.

Procedure - manage Azure AD-based authentication & authorization and region-specific Azure subscriptions

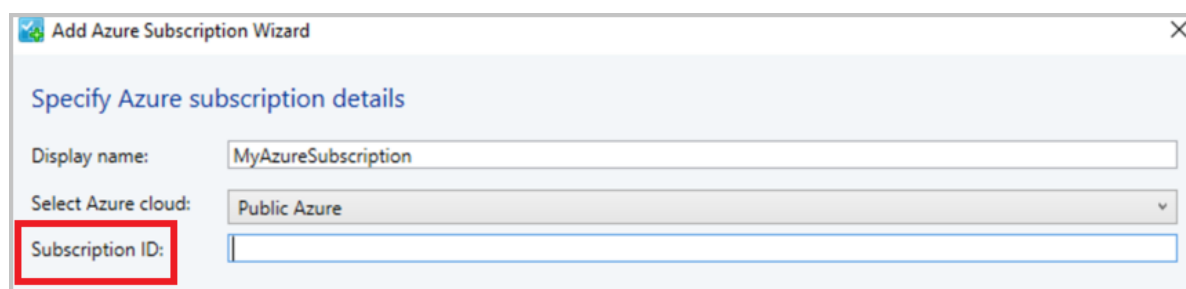
Use the following steps:

1. Browse to **Azure subscription** and click **Add Subscription**.




2. Provide **Display Name**, **Azure cloud**, and **Subscription ID**.

You can provide any friendly name as display name. Choose either public Azure or region-specific subscription as appropriate.



3. Select **Management using Azure AD authentication** (to use certificate based management, go to step 5)



4. Provide **Directory ID**, **Application ID** and **Key**, and then click **Finish** (after this step, directly go to step 6).

