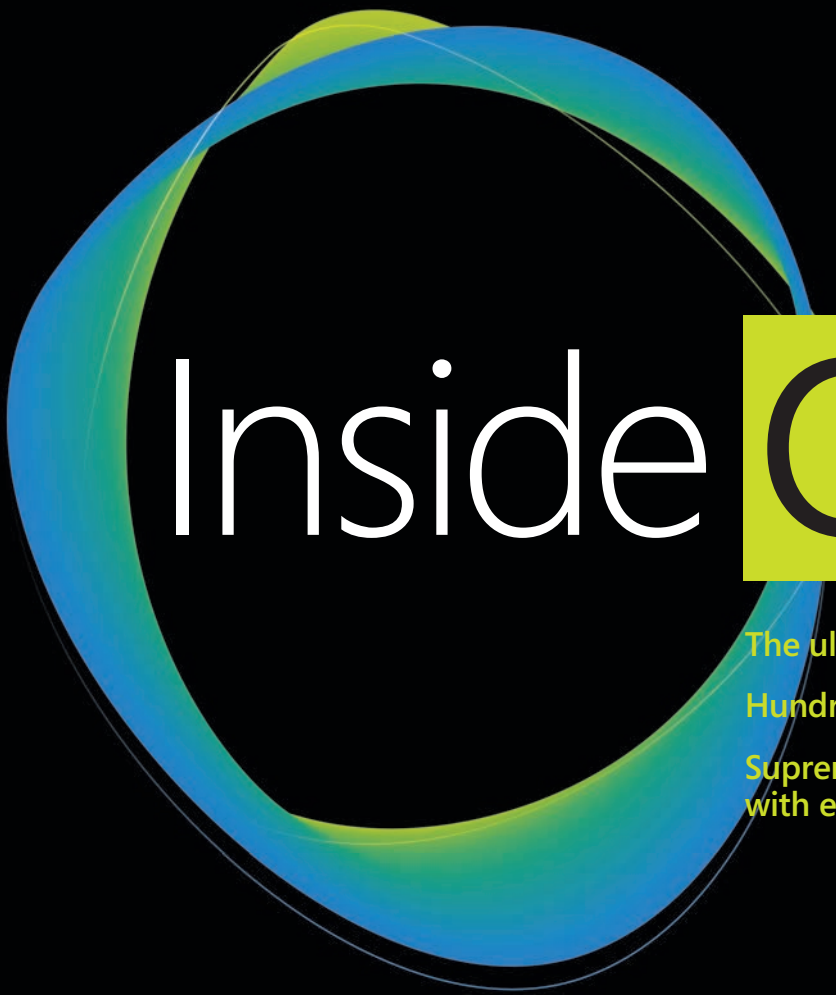

Microsoft

Inside OUT

The ultimate in-depth reference

Hundreds of timesaving solutions

Supremely organized, packed with expert advice

Windows Server 2019

Orin Thomas

Microsoft Cloud Operations Advocate, Cloud and Datacenter expert, and leading Windows author

Microsoft

Windows Server 2019 Inside Out

Orin Thomas

Hyper-V guest cluster storage

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

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

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

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

Shared virtual hard disk

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

Shared virtual hard disks have the following requirements:

- Can be used with generation 1 and generation 2 VMs.
- Can only be used with guest operating systems running Windows Server 2012 or later. If the guest operating systems are running Windows Server 2012, they must be updated to use the Windows Server 2012 R2 integration services components.
- Can only be used if virtualization hosts are running the Windows Server 2012 R2 or later version of Hyper-V.
- Must be configured to use the .vhdx virtual hard disk format.

- Must be connected to a virtual SCSI controller.
- When deployed on a failover cluster, the shared virtual hard disk itself should be located on shared storage, such as a Continuously Available File Share or Cluster Shared Volume. This is not necessary when configuring a guest failover cluster on a single Hyper-V server that is not part of a Hyper-V failover cluster.
- VMs can only use shared virtual hard disks to store data. You can't boot a VM from a shared virtual hard disk.

The configuration of shared virtual hard disks differs from the traditional configuration of VM guest failover clusters because you configure the connection to shared storage by editing the VM properties rather than connecting to the shared storage from within the VM. Windows Server 2019 and Windows Server 2016 support shared virtual hard disks being resized and used with Hyper-V replica.

Hyper-V VHD Sets

VHD Sets are a newer version of shared virtual hard disks. Hyper-V VHD Sets use a new virtual hard disk format that uses the .vhds extension. VHD Sets support online resizing of shared virtual disks, Hyper-V Replicas, and application-consistent Hyper-V checkpoints.

You can create a VHD Set file from Hyper-V Manager or by using the *New-VHD* cmdlet with the file type set to .vhds when specifying the virtual hard disk name. You can use the *Convert-VHD* cmdlet to convert an existing shared virtual hard disk file to a VHD Set file as long as you have taken the VMs that use the shared virtual hard disk file offline and have removed the shared virtual hard disk from the VM using the *Remove-VHHardDiskDrive* cmdlet.

> **More Info**
>
> ***Hyper-V VHD Sets***
>
> **You can learn more about Hyper-V VHD Sets at *https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/manage/create-vhdset-file*.**

Live migration

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

Live migration has the following prerequisites:

- There must be two or more servers running Hyper-V that use processors from the same manufacturer (for example, all Hyper-V virtualization hosts configured with Intel processors or all Hyper-V virtualization hosts configured with AMD processors).
- Hyper-V virtualization hosts need to be members of the same domain, or they must be members of domains that have a trust relationship with each other.
- VMs must be configured to use virtual hard disks or virtual Fibre Channel disks; pass-through disks are not allowed.

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

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

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

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

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

The next step in configuring live migration is choosing which authentication protocol and live migration performance options to use. You select these in the Advanced Features area of the Live Migrations settings. The default authentication protocol is CredSSP (Credential Security Support Provider). CredSSP requires local sign-in to both source and destination Hyper-V virtualization hosts to perform live migration. Kerberos allows you to trigger live migration remotely. To use Kerberos, you must configure the computer accounts for each Hyper-V virtualization

host with constrained delegation for the cifs and Microsoft Virtual System Migration Service services, granting permissions to the virtualization hosts that will participate in the live migration partnership. The performance options allow you to speed up live migration. Compression increases processor utilization. SMB will use SMB Direct if both network adapters used for the live migration process support Remote Direct Memory Access (RDMA) and RDMA capabilities are enabled.

Storage migration

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

For example, you could use storage migration to move VM files from one Cluster Share Volume to another on a Hyper-V failover cluster without interrupting the VM's operation. You have the option of moving all data to a single location, moving VM data to separate locations, or moving only the VM's virtual hard disk. To move the VM's data to different locations, select the items you want to move and the destination locations.

Exporting, importing, and copying VMs

A VM export creates a duplicate of a VM that you can import on the same or different Hyper-V virtualization host. When performing an export, you can choose to export the VM, which includes all its VM checkpoints, or you can choose to export just a single VM checkpoint. Windows Server 2019, Windows Server 2016, and Windows Server 2012 R2 support exporting a running VM. With Hyper-V in Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008, it is necessary to shut down the VM before performing an export.

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

When importing a VM, you can choose from the following options:

- **Register The Virtual Machine In Place (Use The Existing ID).** Use this option when you want to import the VM while keeping the VM files in their current locations. Because this method uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which you wish to import the VM.

- **Restore The Virtual Machine (Use The Existing Unique ID).** Use this option when you want to import the VM while moving the files to a new location; for example, you would choose this option if you are importing a VM that was exported to a network share. Because this method also uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which you wish to import the VM.

- **Copy The Virtual Machine (Create A New Unique ID).** Use this method if you want to create a separate clone of the exported VM. The exported files will be copied to a new location, leaving the original exported files unaltered. A new VM ID is created, meaning that the cloned VM can run concurrently on the same virtualization host as the original progenitor VM. When importing a cloned VM onto the same virtualization host as the original progenitor VM, ensure that you rename the newly imported VM; otherwise, you may confuse the VMs.

VM Network Health Detection

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

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

VM drain on shutdown

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

Domain controller cloning

Windows Server 2019 and Windows Server 2016 support creating copies of domain controllers that are running as VMs as long as certain conditions are met. Cloned domain controllers have the following prerequisites:

- The virtualization host supports VM-GenerationID, a 128-bit random integer that identifies each VM checkpoint. VM-GenerationID is supported by the version of Hyper-V available with Windows Server 2012 and later, as well as some third-party hypervisors.
- The domain controller must be running Windows Server 2012 or later as its operating system.
- The server that hosts PDC emulator Flexible Single Master Operations Role (FSMO) must be contactable. This server must be running Windows Server 2012 or later as its operating system. Remember that you should be running the latest version of Windows Server for domain controllers, so of course, you will be running Windows Server 2019 on the servers hosting your forest and domain FSMO roles!
- The computer account of the domain controller that will serve as the template for cloning must be added to the Cloneable Domain Controllers security group.

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

Shielded virtual machines

Shielded virtual machines are a special type of virtual machine that has a virtual Trusted Platform Module (TPM) chip that is encrypted using BitLocker and can only run on specific approved hosts that support what is known as a guarded fabric. Shielded VMs allow sensitive data and workloads to be run on virtualization hosts without the concern that an attacker or the administrator of the virtualization host might export the virtual machine to gain access to the sensitive data stored there. Only certain operating systems can be used for shielded guest virtual machines. You'll learn more about shielded VMs and guarded virtualization fabrics in Chapter 19, "Hardening Windows Server and Active Directory."