Hyper-V

Windows Server 2012 Hyper-V Component Architecture

Secure multitenancy

In Windows Server® 2012, Hyper-V® provides new security and multitenant isolation capabilities to keep virtual machines isolated—even when they are stored on the same physical server. You can expand this with a fully extensible virtual switch, which enables Microsoft partners to develop plug-ins for enhanced networking and security capabilities. These functionalities provide a solution that addresses the complex security requirements of virtual environments.

Flexible infrastructure

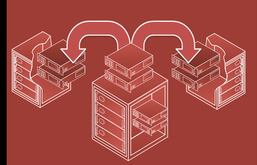

Flexible infrastructure, when and where you need it, is key to easily managing and accessing your virtualized network. With Hyper-V, you can scale beyond virtual local area networks (VLANs) using network virtualization and can place a virtual machine on any node, regardless of its IP address. You can migrate your virtual machines and virtual machine storage flexibly—this includes migrating outside a clustered environment and fully automating management tasks, which reduces the administrative overhead in your environment.

Scalability, performance, and density

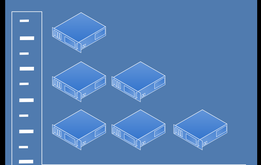

Hyper-V provides support for up to 32 processors and 1 terabyte (TB) of memory for guest operating systems. It also offers a new virtual hard disk format with larger disk capacity that supports up to 64 TB per virtual disk, and provides additional resiliency to enable you to virtualize large-scale workloads. Other new functionalities include resource metering to measure and track consumption of physical resources, support for Offloaded Data Transfer, and improved Quality of Service (QoS) to enforce minimum bandwidth requirements (including network storage requirements).

High availability

Just being able to scale and perform is not enough—you need to ensure that your virtual machines are available when they are needed. Hyper-V provides a wide variety of high availability options. These include simple incremental backup support, enhancements in clustered environments to support up to 4000 virtual machines, parallel live migrations, and encryption with BitLocker® Drive Encryption. You can also use Hyper-V Replica, which replicates virtual machines to an offsite location and provides failover to that site if a failure at the primary site occurs.

Hyper-V Replica

Virtual Machine Replication

Hyper-V Replica is an asynchronous virtual machine replication technology that is included in Windows Server 2012. It is designed for business continuity and disaster recovery. It works with any server, network, or storage vendor. It does not require any shared storage. It enables you to replicate individual or multiple virtual machines. Hyper-V Replica is tightly integrated with Hyper-V and Failover Clustering. You can replicate virtual machines from one computer running Hyper-V at a primary site (the primary server) to another computer running Hyper-V at a Replica site (the Replica server). The Replica server accepts incoming replication traffic from one or more primary servers.

Hyper-V Replica Server Configuration

Configuring a Replica Server

To enable replication, you must configure your Hyper-V server to be a Replica server. This server can accept incoming replication traffic from the virtual machines on your primary (or multiple primary) servers.

Replica Server Authentication

You can set specific authentication types and ports for incoming replication traffic. Options include:

- Kerberos authentication (HTTP) - data sent across the network will not be encrypted.
- Certificate-based authentication (HTTPS) - data sent across the network will be encrypted.

Replica Server Authorization

Hyper-V Replica allows you to specify which servers are permitted to replicate virtual machines to your Replica server.

Replication Storage

Hyper-V Replica allows you to designate a specific location to store your recovery virtual machine files—for example, on a SAN, on an SMB file server, or using direct-attached storage.

Replication Firewall Rules

To allow any incoming virtual machine replication traffic for configured replication ports, you must ensure an inbound firewall rule is created.

Recovery History

Hyper-V Replica frequently replicates changes to your virtual machine, which ensures that your Replica virtual machine is a close match to your primary virtual machine.

Hyper-V Replica can maintain multiple recovery points on the server—these are used to restore a virtual machine. The recovery points contain one or more snapshots. Recovery points are created every hour.

If multiple recovery points are not configured, the server retains only the latest point-in-time recovery point for the Replica virtual machine. You can also choose to take application-consistent snapshots at a specified interval. They use the Volume Shadow Copy Service (VSS).

Maximum Number of Recovery Snapshots

1 — 16

VSS Snapshot Replication Interval

1 Hour — 12 Hours

Hyper-V Replica Failover Operations

Using Hyper-V Replica, you can fail over to a protected virtual machine on the Replica server at any time. There are three main options.

Planned Failover to a Replica Server

A planned failover operation enables you to fail over your production virtual machine to a Replica server as part of a predetermined schedule. Planned failover means no data will be lost. It should be planned for off-business hours.

Two prerequisites are required to perform a planned failover:

- The virtual machine must be shut down before initiating failover.
- The server running Hyper-V at the primary site must be enabled to receive replication traffic. A planned failover also initiates reverse replication.

Test Failover to a Replica Server

You also have the ability to test a Replica virtual machine on the Replica server. The test failover operation does not interrupt ongoing replication. This will create and start a virtual machine with the name "<virtual machine name> - Test".

Unplanned Failover to a Replica Server

In the event of a failure in the primary site, you can bring up the Replica virtual machine on the Replica server. You can choose the most recent replication or an earlier recovery point.

Hyper-V Networking

Load Balancing and Failover

Network adapter teaming, also known as NIC teaming or load balancing and failover (LBFO), enables multiple network adapters to be placed into a team interface. This provides bandwidth aggregation and traffic failover, which prevents loss of connectivity in the event of a network adapter failure. Network adapter teaming supports multivendor implementations.

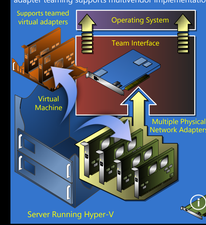

Configurations for Network Adapter Teaming

Switch-independent mode does not require the physical switch to participate in teaming. The switch is not aware that network adapters are part of a team on the server, so the network adapters can be connected to different switches if required.

Switch-dependent modes require a physical switch to participate in teaming. Typically, all the interfaces of the team are connected to the same switch.

Teamed Virtual Network Adapters

Hyper-V supports network adapter teaming within a virtual machine. To provide redundancy, virtual machines should have multiple virtual network adapters, each connected to a different external virtual switch. The virtual machine will have connectivity even if a physical network adapter fails.

Windows Server 2012 supports up to 32 network adapters within a team.

Single Root I/O Virtualization

SR-IOV is a standard that allows PCI Express devices to be shared among multiple virtual machines by providing them a direct hardware path for I/O. Hyper-V provides support for SR-IOV–capable network adapters. SR-IOV reduces network latency, reduces CPU utilization for processing network traffic, and increases network throughput.

SR-IOV–capable networking devices have hardware surfaces called virtual functions that can be securely assigned to virtual machines—bypassing the virtual switch in the management operating system for sending and receiving data. Policy and control remains under the management operating system.

SR-IOV is fully compatible with live migration because software-based networking is available at all times. During live migration, virtual functions are temporarily removed. This enables live migration using network adapters from different vendors, or in a situation where SR-IOV is not available on the destination computer.

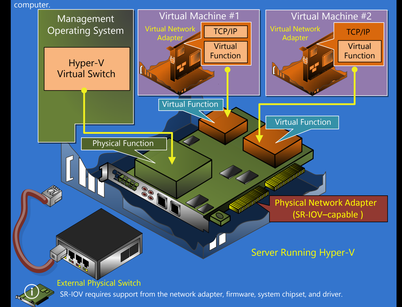

SR-IOV requires support from the network adapter, firmware, system chipset, and driver.

Quality of Service Bandwidth Management

Windows Server 2012 includes new Quality of Service (QoS) bandwidth management functionality that allows you to converge multiple types of network traffic through a single network adapter with a predictable level of service to each type. You can configure bandwidth management features through the virtual machine settings or by using Windows PowerShell commands.

To architect bandwidth management, you can specify a maximum and minimum bandwidth limit. These limits allow you to manage bandwidth allocations depending on your type of network traffic.

It is important to note that the new minimum bandwidth feature allows each network service (such as management, storage, live migration, and virtual machine traffic) to get an allocated share of bandwidth when the network bandwidth is heavily utilized and contended. When bandwidth is freely available, each of these network services gets as much bandwidth as required.

There are two mechanisms to enforce minimum bandwidth. You can use QoS software in your server running Hyper-V, or Windows-certified network adapters that support Data Center Bridging (DCB).

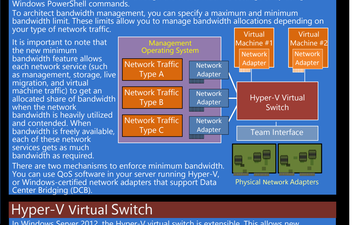

Hyper-V Virtual Switch

In Windows Server 2012, the Hyper-V virtual switch is extensible. This allows new capabilities to be added to the virtual switch so that you can view and manage the traffic on your server running Hyper-V. This includes traffic generated between virtual machines running on the same computer.

Using the extensible capabilities in a virtual switch, Microsoft partners can add their own monitoring, filtering, and forwarding functionality. Any extensions that are created are implemented using Network Driver Interface Specification (NDIS) filter drivers or Windows Filtering Platform (WFP) callout drivers.

You can use NDIS filter drivers to monitor or modify network packets in Windows. You can use WFP to create firewall or intrusion detection functionality.

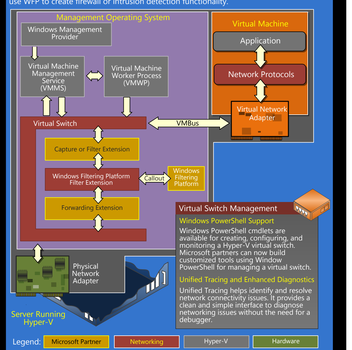

Hyper-V Virtual Machine Mobility

Live Migration Without Shared Storage

Live migration without shared storage (also known as "Shared Nothing Live Migration") enables you to migrate virtual machines and their associated storage between servers running Hyper-V within the same domain. This kind of live migration uses only an Ethernet connection.

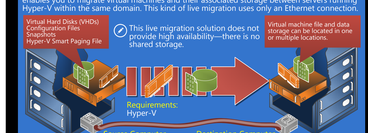

This live migration solution does not provide high availability—there is no shared storage.

Using live migration without shared storage, your virtual machine continues to run while all of its storage is mirrored across to the destination server running Hyper-V. After the Hyper-V storage is synchronized on both sides, the live migration completes its remaining tasks. Finally, the mirror is broken and the source Hyper-V storage is deleted.

- If a failure or a problem occurs during live migration without shared storage, a working virtual machine is always retained.
- Live migration without shared storage can migrate virtual machines between clusters, and from a non-clustered computer to a cluster.
- Live migration without shared storage can migrate virtual machines between different storage types.
- You can also initiate live migration without shared storage using Windows PowerShell.

Storage Migration

Hyper-V storage migration enables you to move virtual machine storage (virtual hard disks) without downtime. This enables new servicing scenarios. For example, you can add more physical disk storage to a non-clustered computer or a Hyper-V cluster and then move the virtual machines to the new storage while the virtual machines continue to run.

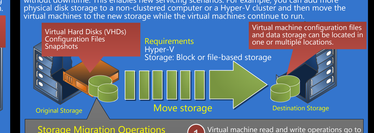

Storage Migration Operations

1. Virtual machine read and write operations go to the VHD on the source device.
2. The VHD is copied from the source VHD storage device to the destination VHD storage device.
3. After the storage is copied, all write operations are mirrored to the source and destination storage devices.
4. After the source and destination storage are synchronized, access to the virtual machine VHD is transferred to the VHD on the destination device.
5. The VHD from the source device is deleted.

You can perform storage migration when the virtual machine is running or when it is turned off. Storage migration moves the storage, not the virtual machine state.

Live Migration with SMB Shared Storage

Live migration with Server Message Block (SMB) shared storage enables you to move virtual machines between servers running Hyper-V within the same domain while the virtual machine storage remains on the SMB-based file server. Concurrent live migrations are supported. This kind of live migration does not require configuration of a failover cluster.

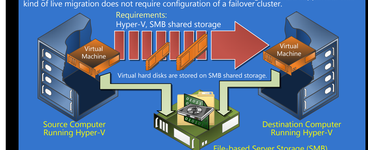

- For this live migration, the virtual hard disk (VHD) resides on an SMB 2.2–based file server. The actual running state of the virtual machine is migrated from one server to another. It is important to note that the connection to the SMB storage is migrated, but the virtual hard disk never moves.
- You can also initiate live migration with SMB shared storage using Windows PowerShell (see the "move-vm" cmdlet).

Live Migration with Failover Clusters

Hyper-V live migration with failover clusters (first introduced in Windows Server 2008 R2) enables you to move running virtual machines from one cluster node running Hyper-V to another node, without any disruption or perceived loss of service. Live migration is initiated by the administrator and is a planned operation.

- In Windows Server 2012, you can select multiple virtual machines within the failover cluster and perform multiple simultaneous live migrations of those virtual machines.
- You can also select and queue live migrations of multiple virtual machines. Live migration queuing is only supported within a failover cluster.

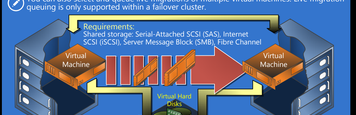

- You can initiate live migration with failover clusters using Windows PowerShell.

Hyper-V Storage

Virtual Fibre Channel for Virtual Machines

Hyper-V virtual Fibre Channel for virtual machines enables virtual machines to access Fibre Channel–based storage. This feature allows you to virtualize workloads that require Fibre Channel storage—and also allows you to cluster guest operating systems in virtual machines using Fibre Channel.

Virtual Fibre Channel Adapters

Virtual Fibre Channel adapters provide port virtualization by exposing host bus adapter (HBA) ports in the guest operating system. This provides the virtual machine with direct and unfiltered access to a storage area network (SAN) using a standard World Wide Name (WWN) that is associated with the virtual machine.

Up to four virtual Fibre Channel adapters can be assigned to each virtual machine.

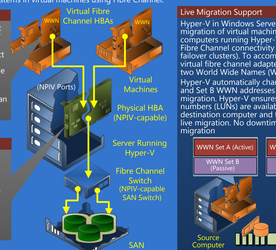

Live Migration Support

Hyper-V in Windows Server 2012 supports live migration of virtual machines across computers running Hyper-V while maintaining Fibre Channel connectivity (live migration with failover clusters). To accomplish this, each virtual fibre channel adapter is configured with two World Wide Names (WWNs).

Hyper-V automatically changes between Set A and Set B WWN addresses during a live migration. Hyper-V ensures that all logical unit numbers (LUNs) are available on the destination computer and then performs the live migration. No downtime occurs during the migration.

WWN Set A (Active)	WWN Set A (Passive)
WWN Set B (Passive)	WWN Set B (Active)

N_Port ID Virtualization (NPIV) Support

Virtual Fibre Channel for virtual machines utilizes N_Port ID Virtualization (NPIV) technology.

An NPIV port is created on the server running Hyper-V and is associated with the virtual Fibre Channel adapter. The World Wide Name (WWN) assigned to the NPIV port allows all I/O to be redirected to a specific virtual Fibre Channel adapter in a virtual machine.

MPIO Connectivity to Fibre Channel Storage

Hyper-V in Windows Server 2012 can use Multipath I/O (MPIO) functionality to ensure continuous connectivity to Fibre Channel storage from within a virtual machine.

MPIO on the Virtual Machine

You can configure multiple virtual Fibre Channel adapters inside a virtual machine, and use a separate copy of Multipath I/O (MPIO) within the guest operating system of the virtual machine to connect to the LUNs that the virtual machine can access.

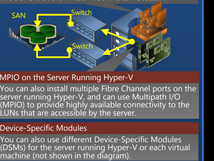

MPIO on the Server Running Hyper-V

You can also install multiple Fibre Channel ports on the server running Hyper-V, and can use Multipath I/O (MPIO) to provide highly available connectivity to the LUNs that are accessible by the server.

Device-Specific Modules

You can also use different Device-Specific Modules (DSMs) for the server running Hyper-V or each virtual machine (not shown in the diagram).

New Virtual Hard Disk Format

VHD is a virtual hard disk file format that enables you to represent a physical hard disk drive in a file, and it is used as the hard disk of a virtual machine. Hyper-V in Windows Server 2012 contains an update to the virtual hard disk format called VHDX.

VHDX Requirements:
Windows Server 2012
Windows 8

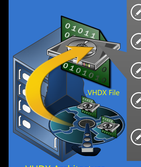

VHDX Format

- Supports up to 64 TB of storage capacity.
- Logs updates to the VHDX metadata structures. This provides added resiliency to the VHDX file in the case of power outages.
- Supports larger block sizes for dynamic and differencing disks, which allows the disks to be tuned to the needs of virtualized workloads.
- Increases performance for applications and workloads, especially on physical disks that have a larger sector size than 512 bytes.
- Supports storing custom metadata. For example, you might want to record your operating system version or any patches you have applied.
- You can also configure and manage virtual hard disks on a computer running Hyper-V using Windows PowerShell commands.

Hyper-V Using Server Message Block (SMB)

Hyper-V can store virtual machine files (configuration files, virtual hard disk files, and snapshots) on file servers using Server Message Block (SMB) 2.2. This is supported for both non-clustered and clustered servers running Hyper-V where file storage is used as shared storage for the failover cluster.

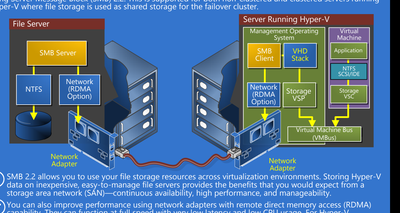

- SMB 2.2 allows you to use file storage resources across virtualization environments. Storing Hyper-V data on inexpensive, easy-to-manage file servers provides the benefits that you would expect from a storage area network (SAN)—continuous availability, high performance, and manageability.
- You can also improve performance using network adapters with remote direct memory access (RDMA) capability. They can function at full speed with very low latency and low CPU usage. For Hyper-V workloads, this allows a remote file server to have performance that compares to local storage.

Hyper-V and Failover Clustering

Clustered Virtual Machines for High Availability

Virtual Machine Live Migrations

Failover Clustering supports concurrent live migrations of virtual machines. The cluster initiates as many live migrations as possible at the same time, and then queues the remaining live migrations.

Using the Failover Cluster Manager, you can monitor the detailed status of on going or queued live migrations.

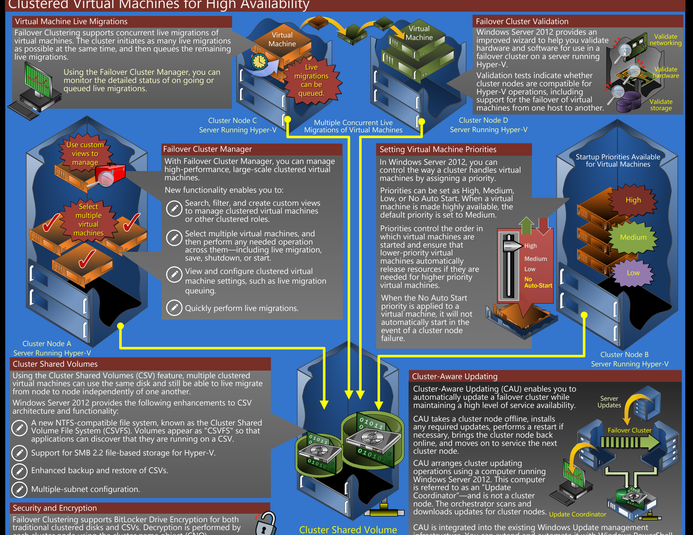

Failover Cluster Validation

Windows Server 2012 provides an improved wizard to help you validate hardware and software for use in a failover cluster on a server running Hyper-V.

Validation tests indicate whether cluster nodes are compatible for Hyper-V operations, including support for the failover of virtual machines from one host to another.

Failover Cluster Manager

With Failover Cluster Manager, you can manage high-performance, large-scale clustered virtual machines.

New functionality enables you to:

- Search, filter, and create custom views to manage clustered virtual machines or other clustered roles.
- Select multiple virtual machines, and then perform any needed operation across them—including live migration, save, shutdown, or start.
- View and configure clustered virtual machine settings, such as live migration queuing.
- Quickly perform live migrations.

Setting Virtual Machine Priorities

In Windows Server 2012, you can control the way a cluster handles virtual machines by assigning a priority.

Priorities can be set as High, Medium, Low, or No Auto Start. When a virtual machine is made highly available, the default priority is set to Medium.

Priorities control the order in which virtual machines are started and ensure that lower-priority virtual machines automatically release resources if they are needed for higher priority virtual machines.

When the No Auto Start priority is applied to a virtual machine, it will not automatically start in the event of a cluster node failure.

Cluster Shared Volumes

Using the Cluster Shared Volumes (CSV) feature, multiple clustered virtual machines can use the same disk and still be able to live migrate from node to node independently of one another.

Windows Server 2012 provides the following enhancements to CSV architecture and functionality:

- A new NTFS-compatible file system, known as the Cluster Shared Volume File System (CSVFS). Volumes appear as "CSVFS" so that applications can discover that they are running on a CSV.
- Support for SMB 2.2 file-based storage for Hyper-V.
- Enhanced backup and restore of CSVs.
- Multiple-subnet configuration.

Security and Encryption

Failover Clustering supports BitLocker Drive Encryption for both traditional clustered disks and CSVs. Decryption is performed by each cluster node using the cluster name object (CNO).

Cluster-Aware Updating

Cluster-Aware Updating (CAU) enables you to automatically update a failover cluster while maintaining a high level of service availability.

CAU takes a cluster node offline, installs any required updates, performs a restart if necessary, brings the cluster node back online, and moves on to service the next cluster node.

CAU arranges cluster updating operations using a computer running Windows Server 2012. This computer is referred to as an "Update Coordinator"—and is not a cluster node. The orchestrator scans and downloads updates for cluster nodes.

CAU is integrated into the existing Windows Update management infrastructure. You can extend and automate it with Windows PowerShell.

Hyper-V Scalability

Physical Hardware and Virtual Machine Scalability

Hyper-V in Windows Server 2012 provides enhanced enterprise hosting capabilities, with expanded support for both physical and virtual processors, and physical and virtual memory. It also makes it easier to virtualize high-performance workloads by supporting the configuration of large, high-performance virtual machines.

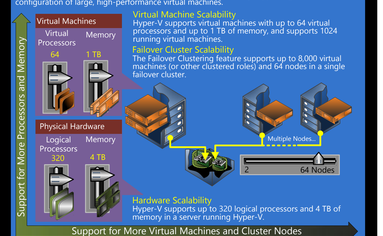

Virtual Machine Scalability

Hyper-V supports virtual machines with up to 64 virtual processors and up to 1 TB of memory, and supports 1024 running virtual machines.

Failover Cluster Scalability

The Failover Clustering feature supports up to 8,000 virtual machines (or other clustered roles) and 64 nodes in a single failover cluster.

Hardware Scalability

Hyper-V supports up to 320 logical processors and 4 TB of memory in a server running Hyper-V.

NUMA and Virtual Machines

NUMA (Non-Uniform Memory Access) is a multiprocessor architecture that groups memory and processors into compute nodes. The time required for a processor to access memory within a node is faster than the time required for a processor to access memory across nodes. Hyper-V supports projecting a virtual NUMA topology within a virtual machine, which enables virtual machines with multiprocessors to scale better.

The guest operating system and applications can take advantage of any NUMA performance optimizations. By default, the virtual NUMA topology within a virtual machine is optimized to match the NUMA topology in the server running Hyper-V.

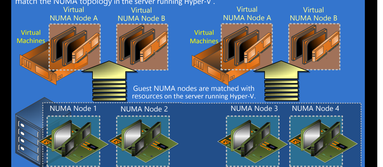

Author: Martin McClean (Server and Cloud Division Information Experience) email: virtua@microsoft.com