

A man in a light blue shirt and dark blue tie is pointing at a whiteboard with a blue marker. The whiteboard has some faint diagrams and text. There are four blue squares on the page: two at the top and two on the man's shirt.

# Windows Server 2012 R2

## Server Virtualization Technical Overview

### Copyright information

© 2013 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other website references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes. You may modify this document for your internal, reference purposes.

# Table of contents

Table of contents .....	1
Beyond Virtualization .....	6
<b>Before Windows Server 2012 R2.....</b>	<b>6</b>
Windows Server 2008 R2 Hyper-V Enhancements .....	7
Windows Server 2008 R2 Hyper-V Benefits.....	7
Windows Server 2012 Hyper V and Windows Server 2012 R2 .....	7
<b>Scalability, Performance &amp; Density .....</b>	<b>8</b>
Virtual Machine NUMA .....	9
Enhanced Storage Capabilities .....	10
Support for Advanced Format Drives (4 KB Sector Disks) in Hyper-V .....	11
New Virtual Hard Disk Format (VHDX) .....	12
Online Virtual Hard Disk Resize .....	15
Online Checkpoint Merge .....	16
Virtual Fibre Channel in Hyper-V .....	17
Offloaded Data Transfer .....	19
Enhanced Networking Performance.....	21
Virtual Receive Side Scaling .....	21
Dynamic Virtual Machine Queue.....	23
Single-Root I/O Virtualization .....	24
Enhanced Resource Management.....	26
Dynamic Memory .....	26
Hyper-V Resource Metering .....	29
Network Quality of Service (QoS) .....	32
Storage Quality of Service (QoS).....	36
Virtualized Workload Performance .....	37

The Hyper-V Extensible Switch .....	42
PVLANs .....	43
ARP/ND Poisoning and Spoofing Protection .....	44
DHCP Guard Protection .....	44
Hyper-V Virtual Switch Extended Port ACLs .....	44
Trunk Mode to Virtual Machines.....	45
Monitoring .....	45
Windows PowerShell and WMI.....	45
Extending the Extensible Switch.....	46
Extensibility .....	47
Manageability .....	49
Partner Extensions.....	49
Physical Security .....	50
BitLocker.....	51

## Flexible Infrastructure..... 53

Linux Support on Hyper-V.....	53
64 vCPU Support & Deeper Integration Services Support .....	54
Synthetic Frame Buffer Driver.....	54
Dynamic Memory Support.....	54
Live Virtual Machine Backup Support .....	54
Dynamic Expansion of Live Fixed Sized VHDXs.....	55
Linux kdump/kexec support.....	55
NMI Support.....	55
Specification of Memory Mapped I/O (MMIO) Gap.....	55
Virtual Machine Live Cloning.....	56
Virtual Machine Mobility .....	57
Live Migration .....	57
SMB-Based Live Migration.....	59
Faster and Simultaneous Migrations .....	59
Live Migration with Compression .....	59
Live Migration over SMB .....	60
Storage Live Migration.....	61
Shared Nothing Live Migration.....	62

Cross-Version Live Migration.....	65
Reliably Import Virtual Machines.....	66
Automation Support for Hyper-V.....	68
Hyper-V Network Virtualization .....	70
Practical Applications.....	72
Network Virtualization - Key Benefits.....	73
Network Virtualization Concepts .....	74
Packet Encapsulation .....	76
Network Virtualization through Address Virtualization .....	77
Network Virtualization Architecture.....	79
Network Virtualization Routing   Between Virtual Subnets.....	80
Network Virtualization Routing   Outside the Virtual Network.....	80
<b>High Availability &amp; Resiliency .....</b>	<b>83</b>
NIC Teaming .....	83
Architecture .....	84
NIC Teaming Configurations .....	84
Traffic Distribution Algorithms.....	85
NIC Teaming In Virtual Machines.....	86
Failover Clustering .....	86
Cluster Shared Volumes.....	87
Active Directory-Detached Clusters .....	89
Cluster Quorum & Dynamic Witness.....	89
VM Drain on Shutdown.....	93
VM Network Health Detection.....	93
Enhanced Cluster Dashboard .....	94
VM Monitoring .....	94
Failover Priority, Affinity and Anti-Affinity.....	96
Cluster-Aware Updating .....	97
Guest Clustering .....	100
Shared VHDX.....	101
Incremental Backup.....	103
Windows Azure Backup Integration.....	104
Hyper-V Replica .....	105

Extended Replication .....	107
Windows Azure Hyper-V Recovery Manager.....	109
<b>Virtualization Innovation .....</b>	<b>111</b>
Generation 2 VMs .....	111
Enhanced Session Mode .....	112
Automatic Virtual Machine Activation.....	114
<b>Conclusion .....</b>	<b>116</b>

# Beyond Virtualization


Server virtualization has evolved over the past few years from a nascent technology into a mature IT feature. In the process, businesses of all shapes and sizes have begun taking advantage of its power to meet shifting business needs. By virtualizing their workloads, organizations can control and cut costs while improving the scalability, flexibility, and reach of IT systems.

With these advances, however, comes the realization that virtualization by itself does not allow organizations to build or take advantage of cloud services, which are assuming an ever-growing role in the execution of business tasks.

Microsoft has taken a leading position in the advancement of virtualization technology with Hyper-V. First introduced as part of Windows Server 2008, and then expanded and enhanced in Windows Server 2008 R2 and again in Windows Server 2012, Hyper-V provides organizations with a tool for optimizing server hardware investments by consolidating multiple server roles as separate virtual machines running on a single physical host machine. They can also use Hyper-V to efficiently run multiple operating systems—including operating systems other than Windows, such as Linux—together on a single server, and take advantage of the power of 64-bit computing.

This whitepaper discusses the significant improvements made to the virtualization capabilities in Windows Server 2012 R2, and how these combine with the existing, powerful capabilities of Windows Server 2012 Hyper-V, to provide customers with a comprehensive platform to handle the demands of the modern datacenter.

## Before Windows Server 2012 R2



Let's first review the Hyper-V improvements that the earlier versions of Windows Server provide. Beginning with Windows Server 2008, in June 2008, server virtualization via Hyper-V technology has been an integral part of the operating system. A new version of Hyper-V was included as a part of Windows Server 2008 R2, and this was further enhanced with Service Pack 1 (SP1).

There are two manifestations of the Hyper-V technology:

- **Hyper-V** is the hypervisor-based virtualization role of **Windows Server**.

- **Microsoft Hyper-V Server** is the hypervisor-based server virtualization product that allows customers to consolidate workloads onto a single physical server. This is available as a free download.

## Windows Server 2008 R2 Hyper-V Enhancements

With the launch of Windows Server 2008 R2 Hyper-V, in October 2009, Microsoft introduced a number of compelling capabilities to help organizations reduce costs, whilst increasing agility and flexibility. Key features introduced included:

- **Live Migration** – Enabling the movement of virtual machines (VMs) with no interruption or downtime
- **Cluster Shared Volumes** – Highly scalable and flexible use of shared storage (SAN) for VMs
- **Processor Compatibility** – Increase the Flexibility for Live Migration across hosts with differing CPU architectures
- **Hot Add Storage** – Flexibly add or remove storage to and from VMs
- **Improved Virtual Networking Performance** – Support for Jumbo Frames and Virtual Machine Queue (VMq)

With the addition of Service Pack 1 (SP1) for Hyper-V, in October 2011, Microsoft introduced 2 new, key capabilities to help organizations realize even greater value from the platform:

- **Dynamic Memory** – More efficient use of memory while maintaining consistent workload performance and scalability.
- **RemoteFX** – Provides the richest virtualized Windows 7 experience for Virtual Desktop Infrastructure (VDI) deployments.

## Windows Server 2008 R2 Hyper-V Benefits

Hyper-V is an integral part of Windows Server and provides a foundational virtualization platform that lets customers transition to the cloud. With Windows Server 2008 R2, customers get a compelling solution for core virtualization scenarios; production server consolidation, dynamic data center, business continuity, Virtual Desktop Infrastructure (VDI), and test and development. Hyper-V provides customers with better flexibility with features like live migration and cluster shared volumes for storage flexibility. In Windows Server 2008 R2, Hyper-V also delivered greater scalability with support for up to 64 logical processors and improved performance with support for dynamic memory and enhanced networking support.

## Windows Server 2012 Hyper V and Windows Server 2012 R2

Fast forward to September 2012, and the launch of Windows Server 2012. This brought an incredible number of new and an enhanced Hyper-V Capabilities. These capabilities, many of which we'll discuss in this paper, ranged from enhancements around scalability, new storage and networking features, significant enhancements to the Live Migration capabilities, deeper integration with hardware, and an in-box VM replication capability, to name but a few. These improvements, new features and enhancements can be grouped into 4 key areas, and it's these key areas we'll focus on throughout this whitepaper, looking at both Windows Server 2012 and R2. The 4 key areas are:

- **Scalability, Performance & Density** – customers are looking to run bigger, more powerful virtual machines, to handle the demands of their biggest workloads. In addition, as hardware scale grows, customers wish to take advantage of the largest physical systems to drive the highest levels of density, and reduce overall costs.

- **Security & Multitenancy** - Virtualized data centers are becoming more popular and practical every day. IT organizations and hosting providers have begun offering infrastructure as a service (IaaS), which provides more flexible, virtualized infrastructures to customers—"server instances on-demand." Because of this trend, IT organizations and hosting providers must offer customers enhanced security and isolation from one another, and in some cases, encrypted to meet compliance demands.
- **Flexible Infrastructure** – In a modern datacenter, customers are looking to be agile, in order to respond to changing business demands quickly, and efficiently. Being able to move workloads flexibly around the infrastructure is of incredible importance, and in addition, customers want to be able to choose where best to deploy their workloads based on the needs of that workload specifically.
- **High Availability & Resiliency** – As customers' confidence in virtualization grows, and they virtualize their more mission-critical workloads, the importance of keeping those workloads continuously available grows significantly. Having capabilities built into the platform that not only help keep those workloads highly available, but also, in the event of a disaster, quick to restore in another geographical location, is of immense importance when choosing a platform for today's modern datacenter.

Throughout these 4 areas, we'll explore the challenges that customers face, and how the capabilities found within Windows Server 2012 R2 can help address those challenges with powerful, yet cost effective solutions.

## Scalability, Performance & Density

Hyper-V in Windows Server 2008 R2 supported configuring virtual machines with a maximum of four virtual processors and up to 64 GB of memory. However, IT organizations increasingly want to use virtualization when they deploy mission-critical, tier-1 business applications. Large, demanding workloads such as online transaction processing (OLTP) databases and online transaction analysis (OLTA) solutions typically run on systems with 16 or more processors and demand large amounts of memory. For this class of workloads, more virtual processors and larger amounts of virtual machine memory are a core requirement.

Scalability however, goes beyond just running workloads. Customers also need to ensure that the demands of workloads can be handled effectively by scalable storage and networking infrastructure, and to do so, must take advantage of the latest, and greatest hardware innovations.

With Windows Server 2012, and subsequently 2012 R2, there were a number of design goals to try to address these challenges. Not only do we want to enable customers to run their most demanding of applications, whilst providing the highest levels of performance and scale, but at the same time, we want to ensure that customers can provide optimal resource usage and availability across their infrastructure.

From an out and out scalability perspective, Hyper-V in Windows Server 2012 R2 greatly expands support for host processors and memory over Windows Server 2008 R2 Hyper-V. New features include support for up to 64 virtual processors and 1TB of memory for Hyper-V guests, a new VHDX virtual hard disk format with larger disk capacity of up to 64 TB, and additional resiliency and alignment optimization, which we'll discuss later. These features help ensure that the virtualization infrastructure can support the configuration of large, high-performance virtual machines to support workloads that might need to scale up significantly.

These however, aren't the only improvements in Windows Server 2012 Hyper-V, as you can see from the table below:



	Resource	Windows Server 2008 R2 Hyper-V	Windows Server 2012 R2 Hyper-V	Improvement Factor
<b>Host</b>	Logical Processors	64	<b>320</b>	<b>5x</b>
	Physical Memory	1TB	<b>4TB</b>	<b>4x</b>
	Virtual CPUs per Host	512	<b>2,048</b>	<b>4x</b>
<b>VM</b>	Virtual CPUs per VM	4	<b>64</b>	<b>16x</b>
	Memory per VM	64GB	<b>1TB</b>	<b>16x</b>
	Active VMs per Host	384	<b>1,024</b>	<b>2.7x</b>
	Guest NUMA	No	<b>Yes</b>	-
<b>Cluster</b>	Maximum Nodes	16	<b>64</b>	<b>4x</b>
	Maximum VMs	1,000	<b>8,000</b>	<b>8x</b>

Table 1 – Comparison of Windows Server 2008 R2 Hyper-V & Windows Server 2012 R2 Hyper-V Scalability

From a host perspective, you can see from the table that Hyper-V supports up to 4TB of physical memory per host, and up to 2,048 vCPUs per host. This is a 4x increase over Windows Server 2008 R2 Hyper-V, and means that a customer could, in reality, run 1,024 2-vCPU virtual machines, each with around 4GB memory, and still be within a supported configuration. This scalability is immense, and ensures customers can realize the greatest value for their hardware investments.

When we think about Virtual Machines (VM) in particular, again, significant improvements have been made across the board, with Hyper-V now supporting VMs with up to 64 vCPUs, and 1TB memory. This is huge scale, and opens the door to running high-end, mission-critical in-memory transactional or analysis workloads that can benefit significantly from that kind of resource capacity.

Earlier, we briefly discussed how customers are demanding higher levels of availability and resiliency for their key virtualized workloads. With Windows Server and Hyper-V, the foundation of providing that higher level of availability is the Failover Cluster. With Windows Server 2012 R2, cluster sizes have increased from a maximum of 16 nodes in Windows Server 2008 R2, to 64 nodes in Windows Server 2012 and Windows Server 2012 R2. This in turn, supports a significantly higher number of active virtual machines per cluster, up from 1,000 to 8,000.

There is one other innovation, highlighted in the table that can drive higher levels of performance for virtualized workloads, and is of particular importance when running virtualized workloads with significant numbers of virtual processors, and high levels of memory. That innovation is Virtual Machine NUMA.

## Virtual Machine NUMA

Windows Server 2012 R2 Hyper-V now supports NUMA, or Non-Uniform Memory Access, inside a virtual machine. NUMA refers to a computer architecture in multiprocessor systems, in which the required time for a processor to access memory depends on the memory's location relative to the processor

With NUMA, a processor can access local memory (memory attached directly to the processor) faster than it can access remote memory (memory that is local to another processor in the system). Modern operating systems and high-performance applications such as SQL Server have developed optimizations to recognize the system's NUMA topology and consider NUMA when they schedule threads or allocate memory to increase performance.

Projecting a virtual NUMA topology into a virtual machine provides optimal performance and workload scalability in large virtual machine configurations. It does this by letting the guest operating system and applications such as SQL Server, or the Windows Web Server, IIS, take advantage of their inherent NUMA performance optimizations. The default virtual NUMA topology that is projected into a Hyper-V virtual machine is optimized to match the host's NUMA topology, as shown in the following figure.

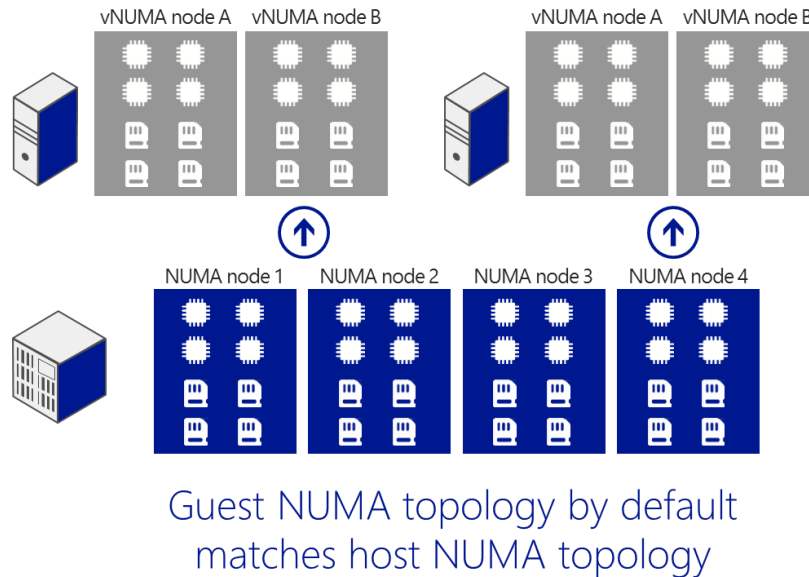


Figure 1 – Virtual Machine NUMA nodes aligning with physical NUMA topology

### Why This Matters

**Guest NUMA ensures that key workloads that have NUMA-aware capabilities, can perform at their highest possible levels, and take advantage of the underlying performance characteristics and capabilities of the hardware itself, maximizing the investment in both hardware, software and the applications. Customers running SQL, and IIS will benefit considerably from Guest NUMA.**

**High-performance applications such as Microsoft SQL Server 2012 and Internet Information Services (IIS) 8 in Windows Server 2012 are NUMA-aware, enabling significant performance enhancements over virtualized instances of the application on non-NUMA aware platforms and VMs. Guest NUMA support also works for high-availability solutions using Windows Server 2012 failover clustering. Failover clusters evaluate the NUMA configuration of a node before moving a VM; this ensures that the target node is able to support the VM's workload.**

## Enhanced Storage Capabilities

Windows Server 2012 introduced a number of new and powerful storage capabilities that can play a significant role in a virtualized infrastructure, in order to support the most intensive, mission-critical of workloads. With Windows Server 2012 R2, there have been further enhancements that increase performance and flexibility and help to ensure continuous availability.

## Support for Advanced Format Drives (4 KB Sector Disks) in Hyper-V

Increases in storage density and reliability are among the factors driving the data storage industry to transition the physical format of hard disk drives from 512-byte sectors to 4,096-byte sectors (also known as 4 KB sectors). However, most of the software industry depends on disk sectors of 512 bytes in length. A change in sector size introduces major compatibility issues in many applications. To minimize the impact on the ecosystem, hard drive vendors are introducing transitional "512-byte emulation drives" also known as "512e." These drives offer some advantages of 4 KB native drives, such as improved format efficiency and an improved scheme for error correction codes (ECC), but with fewer compatibility issues than by exposing a 4 KB sector size at the disk interface. Hyper-V in Windows Server 2012 and Windows Server 2012 R2 supports "512e" and 4 KB disk sectors.

Customers face a challenge of ensuring they can adopt and take advantage of this emerging disk format to provide the best performance and optimization for their key workloads.

Support for 4,096-byte sectors (4 KB disk sectors) in virtual disks, a standard to which the industry will move over the next few years to support increasing storage requirements, was first introduced in Windows Server 2012 Hyper-V. Hyper-V in Windows Server 2012, and subsequently Windows Server 2012 R2 Hyper-V, also provides enhanced performance of the transitional standard, 512-byte emulation drives, also known as 512-byte Emulation (512e). Support for 4 KB disk sectors and 512e helps ensure that your virtualization infrastructure keeps pace with industry innovations in storage.

### Hyper-V & 512e disks

A 512e disk can perform a write only in terms of a physical sector, that is, it cannot directly write a 512-byte sector write that is issued to it. The internal process in the disk which makes this write possible follows these steps:

1. The disk reads the 4 KB physical sector into its internal cache, which contains the 512-byte logical sector that is referred to in the write.
2. Data in the 4 KB buffer is modified to include the updated 512-byte sector.
3. The disk performs a write of the updated 4 KB buffer back to its physical sector on the disk.

This above process is called "Read-Modify-Write," or RMW. The RMW process causes performance degradation in virtual hard disks for the following reasons:

- Dynamic and differencing virtual hard disks have a 512-byte sector bitmap in front of their data payload. In addition, footer/header/parent locators all align to a 512-byte sector. It is common for the virtual hard disk drive to issue 512-byte writes to update these structures, resulting in the RMW behavior described earlier.
- Applications commonly issue reads and writes in multiples of 4 KB sizes (the default cluster size of NTFS). Because a 512-byte sector bitmap is in front of the data payload block of dynamic and differencing virtual hard disks, the 4 KB blocks are not aligned to the physical 4 KB boundary, as shown in the following figure.

In the figure below, the virtual hard disk 4 KB block is not aligned with physical 4 KB boundary




Figure 2 - Virtual hard disk 4 KB block (blue) not aligned with physical 4 KB boundary

Each 4 KB write issued by the legacy parser, to update the payload data, results in two reads for two blocks on the disk, which are then updated and subsequently written back to the two disk blocks. The overall performance impact of the RMW process on workloads usually ranged 30 to 80 percent and, at times, was even higher.

The overall performance impact of the RMW process on workloads usually ranged 30 to 80 percent and, at times, was even higher.

Hyper-V in Windows Server 2012 R2 mitigates the performance-degrading effect of 512e disks on the virtual hard disk stack by preparing the previously mentioned structures for alignment to 4 KB boundaries in the VHD format. This avoids the RMW effect when accessing data within the virtual hard disk file and when updating virtual hard disk metadata structures.

### Native 4KB Sector Support

Hyper-V in Windows Server 2012 R2 makes it possible to store virtual hard disks on 4 KB disks by implementing a software RMW algorithm in the virtual hard disk layer. This algorithm converts 512-byte access-and-update requests to corresponding 4 KB accesses and updates.

### Requirements

To take advantage of Hyper-V support for 4 KB disk sectors, you need the following:

- Windows Server 2012 with Hyper-V, **Windows Server 2012 R2 with Hyper-V**, Hyper-V Server 2012 or **Hyper-V Server 2012 R2**
- Physical disk drives that use the 512e or the native 4 KB format.

### Why This Matters

**With the introduction of larger VHDX files (one VHDX disk supports up to 64 terabytes of storage) and ReFS (Resilient File System) volumes in Windows Server 2012 R2, support for 4K sector disks was critical in delivering the capacity and scaling needed to take on the every growing storage needs of customers, without sacrificing performance.**

### New Virtual Hard Disk Format (VHDX)

With the evolution of storage systems, and the ever-increasing reliance on virtualized enterprise workloads, the VHD format of Windows Server needed to also evolve. The new format is better suited to address current and future requirements for running enterprise-class workloads, specifically:

- Where the size of the VHD is larger than 2 TB.
- To reliably protect against issues for dynamic and differencing disks during power failures.

- To prevent performance degradation issues on the new, large-sector physical disks.

Windows Server 2012 Hyper-V introduced a significant update to the VHD format, called VHDX, which has much larger capacity and additional resiliency. **VHDX supports up to 64 terabytes of storage.** It also provides additional protection against corruption from power failures by logging updates to the VHDX metadata structures, and it prevents performance degradation on large-sector physical disks by optimizing structure alignment.

There are a number of new capabilities that are provided by the new VHDX format:

- **Capacity** – support for up to 64TB per virtual disk, and each Hyper-V virtual machine can support up to 256 virtual disks, for a total of petabytes of storage.
- **Corruption Protection** - Protection against corruption during power failures by logging updates to the VHDX metadata structures. The format contains an internal log that is used to capture updates to the metadata of the virtual hard disk file before being written to its final location. In case of a power failure, if the write to the final destination is corrupted, then it is played back from the log to promote consistency of the virtual hard disk file.
- **Optimal Structure Alignment** – Alignment to suit large sector disks. If unaligned I/O's are issued to these disks, an associated performance penalty is caused by the Read-Modify-Write cycles that are required to satisfy these I/O's. The structures in the format are aligned to help ensure that are no unaligned I/O's exist.

There are also a number of other features that are unlocked through the use of the VHDX format.

- **Larger block sizes for dynamic and differential disks** - lets these disks attune to the needs of the workload
- A 4-KB logical sector virtual disk that results in **increased performance** when applications and workloads that are designed for 4-KB sectors use it
- The ability to **store custom metadata** about the file that you might want to record, such as operating system version or patches applied
- Efficiency (called **trim**) in representing data, which results in smaller files and lets the underlying physical storage device reclaim unused space. Trim requires pass-through or SCSI disks and trim-compatible hardware.

The figure below illustrates the VHDX hard disk format.

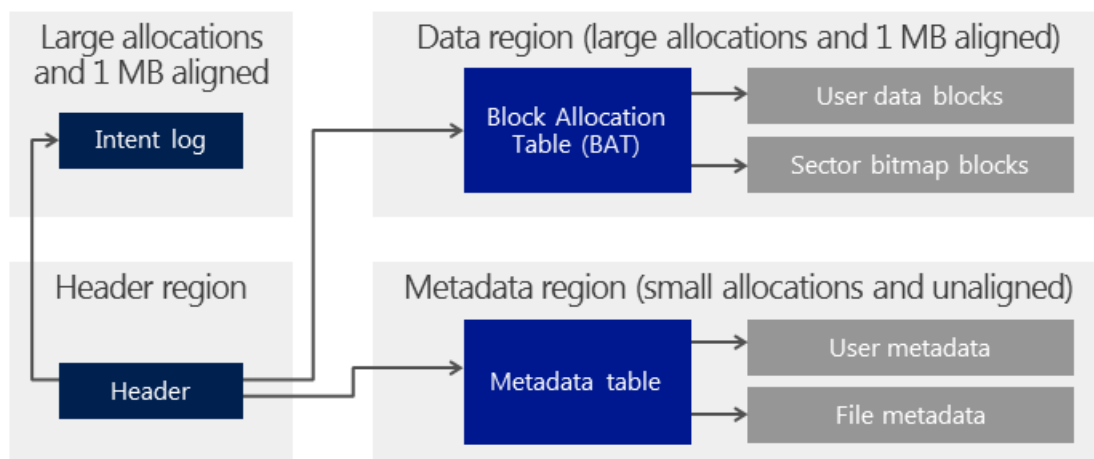


Figure 3 - The VHDX hard disk format

As you can see in the preceding figure, most of the structures are large allocations and are MB aligned. This alleviates the alignment issue that is associated with virtual hard disks. The different regions of the VHDX format are as follows:

- **Header region** - The header region is the first region of the file and identifies the location of the other structures, including the log, block allocation table (BAT), and metadata region. The header region contains two headers, only one of which is active at a time, to increase resiliency to corruptions.
- **Intent log** - The intent log is a circular ring buffer. Changes to the VHDX metastructures are written to the log before they are written to the final location. If corruption occurs during a power failure while an update is being written to the actual location, on the subsequent open, the change is applied again from the log, and the VHDX file is brought back to a consistent state. The log does not track changes to the payload blocks, so it does not protect data contained within them.
- **Data region** - The BAT contains entries that point to both the user data blocks and sector bitmap block locations within the VHDX file. This is an important difference from the VHD format because sector bitmaps are aggregated into their own blocks instead of being appended in front of each payload block.
- **Metadata region** - The metadata region contains a table that points to both user-defined metadata and virtual hard disk file metadata such as block size, physical sector size, and logical sector size.

### VHDX: Better Efficiency & Resiliency

Hyper V in Windows Server 2012, and subsequently, Windows Server 2012 R2, also introduces support that lets VHDX files be more efficient when they represent that data within it.

As the VHDX files can be large, based on the workload they are supporting, the space they consume can grow quickly. Currently, when applications delete content within a virtual hard disk, the Windows storage stack in both the guest operating system and the Hyper V host have limitations that prevent this information from being communicated to the virtual hard disk and the physical storage device.

This contains the Hyper-V storage stack from optimizing the space used and prevents the underlying storage device from reclaiming the space previously occupied by the deleted data.

In Windows Server 2012 R2 Hyper V supports **unmap notifications**, which enables VHDX files be more efficient in representing that data within them. This results in smaller files size, which lets the underlying physical storage device reclaim unused space.

### Requirements

To take advantage of the new VHDX format, you need the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or **Hyper-V Server 2012 R2**

To take advantage of the trim feature, you need the following:

- VHDX-based virtual disks connected as virtual SCSI devices or as directly attached physical disks (sometimes referred to as pass-through disks). This optimization also is supported for natively attached VHDX-based virtual disks.
- Trim-capable hardware

### Why This Matters

**Designed to handle current and future workloads, VHDX has a much larger storage capacity than earlier VHD formats to address the technological demands of evolving enterprises. Performance-enhancing features in VHDX make it easier to handle large workloads, protect data better during power outages, and optimize structure alignments of dynamic and differential disks to prevent performance degradation on new, large-sector physical disks.**

## Online Virtual Hard Disk Resize

Windows Server 2012 R2 introduces a significant improvement within Hyper-V that allows the virtual hard disk file to be resized as needed whilst the VM is running. You can't always predict when a virtual disk will need to be expanded due to new demands in capacity, but equally important is the ability to reclaim space if no longer required. Prior to Windows Server 2012 R2, a VM had to be shut down in order to expand or shrink the virtual hard disk files. Now with Windows Server 2012 R2, this is an online operation, with no downtime for the workload itself. The obvious benefit to this is increased availability and better SLA compliancy.

### Virtual Disk Expansion

Customers have the flexibility to grow the **VHD** or **VHDX** files that are attached to running virtual machines. The administrator will first expand the virtual disk using the Hyper-V management tools, or PowerShell, and then within the Guest OS, the administrator can expand the volume within the OS, using Disk Manager. An example command to resize using PowerShell would be as follows:

```
PS C:\> Resize-VirtualDisk -FriendlyName "Sqldata5" -Size (25GB)
```

Note, this command could also be used for shrinking the virtual disk.

### Virtual Disk Shrinking

Customers also have the flexibility to shrink the **VHDX** files that are attached to running virtual machines. The administrator will first shrink the volume within the Guest OS, and then from within the Hyper-V Management tools, or using PowerShell, shrink the virtual disk. The shrink size will match the space that is freed up inside the VM volume. Note, **Shrink is limited to VHDX only.**

### Requirements

To take advantage of the online adjustment of virtual hard disk size you will need:

- Windows Server 2012 R2 Hyper-V or Hyper-V Server 2012 R2
- VHDX files attached to the Virtual SCSI Controller.

## Why This Matters

**The ability to grow, and shrink virtual disks whilst the virtual machine is running provides customers with significant advantages in flexibility, and provides fewer reasons for needing to introduce downtime for that workload for capacity reasons. Not only does this allow virtual machines to grow flexibly, as data consumption increases inside the VM, but it also allows IT admins to reclaim wasted space that has been allocated, but not used, within a VM. This could be of particular importance in a Service Provider, or Hosting scenario, where a customer is paying for a 100GB VM, but only uses 30GB, and wishes to reduce the disk size to**

**50GB, to reduce their costs. Now this operation can be performed online, with no downtime to the workload.**

## Online Checkpoint Merge

Checkpoints (formerly snapshots) have been mainly used for testing changes to existing virtual machine environments, as a way to return to a previous state or time if required. Having an easier way to revert a virtual machine can be very useful if you need to recreate a specific state or condition so that you can troubleshoot a problem.

Under certain circumstances it makes sense to use checkpoints in a production environment. For example, you can use checkpoints to provide a way to revert a potentially risky operation in a production environment, such as applying an update to the software running in the virtual machine. After successfully testing new changes or updates, many organizations merge their checkpoints back into the original parent disk (to reduce storage space and increase virtual machine disk performance). However, this operation would pause the running virtual machine, effectively making it unavailable while the merge takes place.

In Windows Server 2012 R2 the Hyper-V Live Merge feature now allows organizations to merge current checkpoints back into the original parent while the virtual machine continues to run.

The Hyper-V virtual machine checkpoint feature provides a fast and straightforward way to revert the virtual machine to a previous state. Checkpoint data files (the current leaf node of virtual hard disk that are being forked into a read-only parent differential disk) are stored as .avhd files. **When a checkpoint is deleted, the associated .avhd disks cannot be removed while the virtual machine is running. Windows Server 2012 R2 provides the ability to merge the associated .avhd disk into the parent while the virtual machine continues to run.**

As the process proceeds, I/O is suspended to a small range while data in that range is read from the source and written to the destination. When the leaf is being merged away, further writes to areas that have already been merged are redirected to the merge destination. Upon completion, the online merge fixes the running chain to unlink merged disks and closes those files.

### Requirements

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2

### Why This Matters

**Virtual machine checkpoints capture the state, data, and hardware configuration of a running virtual machine. Many organizations use checkpoints in their current environments for testing updates and patches. However, merging a checkpoint into the parent virtual machine requires downtime and virtual machine unavailability. Now, with the **Live Merge feature of Windows Server 2012 R2 Hyper-V**, you can merge checkpoints into the virtual machine parent while the server is running, with little effect on users. Live merging of checkpoints provides a faster, easier way to revert a virtual machine to a previous state.**



## Virtual Fibre Channel in Hyper-V

Many enterprises have already invested in Fibre Channel SANs, deploying them within their datacenters to address growing storage requirements. These customers often want the ability to utilize this storage from within their virtual machines instead of having the storage accessible to and used only by the Hyper-V host. In addition, customers are looking to achieve true SAN line speed from the VMs, to the SAN.

### Unmediated SAN Access

Virtual Fibre Channel for Hyper-V provides the guest operating system with unmediated access to a SAN by using a standard World Wide Name (WWN) that is associated with a virtual machine. Hyper-V lets you use Fibre Channel SANs to virtualize workloads that require direct access to SAN logical unit numbers (LUNs). Fibre Channel SANs also let you operate in new scenarios, such as running the Windows Failover Clustering feature inside the guest operating system of a virtual machine connected to shared Fibre Channel storage.

### A Hardware-Based I/O Path to the Windows Software Virtual Hard Disk Stack

Mid-range and high-end storage arrays include advanced storage functionality that helps offload certain management tasks from the hosts to the SANs. Virtual Fibre Channel offers an alternative, hardware-based I/O path to the Windows software virtual hard disk stack. This path lets you use the advanced functionality of your SANs directly from within Hyper-V virtual machines. For example, Hyper-V users can offload storage functionality (such as taking a snapshot of a LUN) to the SAN hardware simply by using a hardware Volume Shadow Copy Service (VSS) provider from within a Hyper-V virtual machine.

### Live Migration Support

To support live migration of virtual machines across Hyper-V hosts while maintaining Fibre Channel connectivity, two WWNs, Set A and Set B, are configured for each virtual Fibre Channel adapter. Hyper-V automatically alternates between the Set A and Set B WWN addresses during live migration. This helps ensure that all LUNs are available on the destination host before the migration and that no downtime occurs during the migration. The live migration process that maintains Fibre Channel connectivity is illustrated in the following figure:

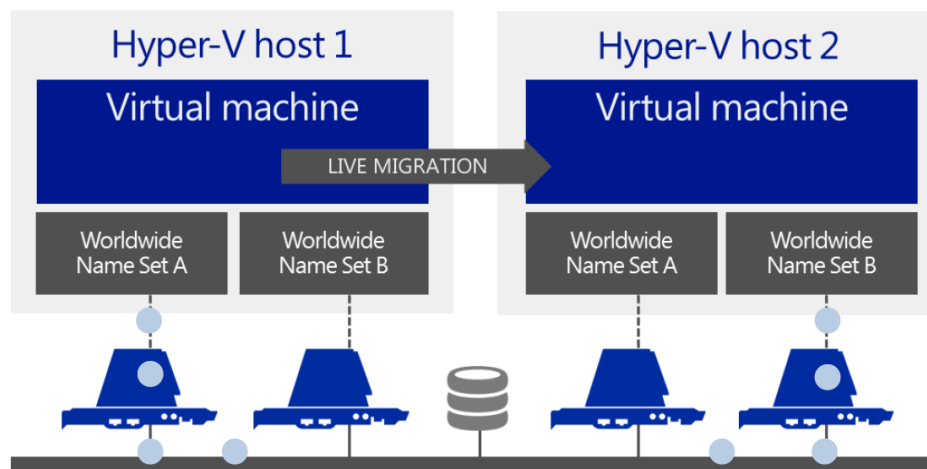


Figure 4 - Alternating WWN addresses during a live migration

### N\_Port ID Virtualization (NPIV)

NPIV is a Fibre Channel facility that lets multiple N\_Port IDs share a single physical N\_Port. This lets multiple Fibre Channel initiators occupy a single physical port, easing hardware requirements in SAN design, especially where virtual SANs are called for. Virtual Fibre Channel for Hyper-V guests uses NPIV (T11 standard) to create multiple NPIV ports on top of the host's physical Fibre Channel ports. A new NPIV port is created on the host each time a virtual HBA is created inside a virtual machine. When the virtual machine stops running on the host, the NPIV port is removed.

### **Flexible Host to SAN Connectivity**

Hyper-V lets you define virtual SANs on the host to accommodate scenarios in which a single Hyper-V host is connected to different SANs via multiple Fibre Channel ports. A virtual SAN defines a named group of physical Fibre Channel ports that are connected to the same physical SAN. For example, assume that a Hyper-V host is connected to two SANs—a production SAN and a test SAN. The host is connected to each SAN through two physical Fibre Channel ports. In this example, you might configure two virtual SANs—one named "Production SAN" that has the two physical Fibre Channel ports connected to the production SAN and one named "Test SAN" that has two physical Fibre Channel ports connected to the test SAN. You can use the same technique to name two separate paths to a single storage target.

### **4 vFC Adapters per VM**

You can configure as many as four virtual Fibre Channel adapters on a virtual machine, and associate each one with a virtual SAN. Each virtual Fibre Channel adapter is associated with one WWN address, or two WWN addresses to support live migration. Each WWN address can be set automatically or manually.

### **Multipath I/O (MPIO)**

Hyper-V in Windows Server 2012 R2 uses Multipath I/O (MPIO) functionality to help ensure optimal connectivity to Fibre Channel storage from within a virtual machine. You can use MPIO functionality with Fibre Channel in the following ways:

- Virtualize workloads that use MPIO. Install multiple Fibre Channel ports in a virtual machine, and use MPIO to provide highly available connectivity to the LUNs that the host can access.
- Configure multiple virtual Fibre Channel adapters inside a virtual machine, and use a separate copy of MPIO within the guest operating system of the virtual machine to connect to the LUNs that the virtual machine can access. This configuration can coexist with a host MPIO setup.
- Use different device specific modules (DSMs) for the host or each virtual machine. This approach permits migration of the virtual machine configuration, including the configuration of DSM and connectivity between hosts and compatibility with existing server configurations and DSMs.

### **Requirements**

Virtual Fibre Channel support in Hyper-V requires the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2
- A computer with one or more Fibre Channel HBAs, each with an updated HBA driver that supports Virtual Fibre Channel. Check with your HBA vendor for information on whether your HBA supports Virtual Fibre Channel.
- Virtual machines that are configured to use a virtual Fibre Channel adapter, which must use Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 as the guest operating system.
- Connection only to data LUNs. Storage accessed through a virtual Fibre Channel that is connected to a LUN cannot be used as restart media.

## Why This Matters

**Virtual Fibre Channel lets you access and use Fibre Channel SAN data from your virtual machines instead of having the storage accessible and used only by the Hyper-V host. Support for Fibre Channel in Hyper-V guests also includes support for many related features, such as NPIV, virtual SANs, live migration, and MPIO. This feature protects your investments in Fibre Channel, lets you virtualize workloads that use direct access to Fibre Channel storage, lets you cluster guest operating systems over Fibre Channel, and offers an important new storage option for servers that are hosted on your virtualization infrastructure.**

## Offloaded Data Transfer

Offloaded Data Transfer (ODX) provides CPU and network offloading to SAN hardware during file copy and move operation between SAN drives. It is especially significant in the cloud space when you must provision new virtual machines from virtual machine template libraries or when virtual hard disk operations are triggered and require large blocks of data to be copied, as in virtual hard disk merges, storage migration, and live migration. These copy operations are then handled by the storage device that must be able to perform offloads (such as an offload-capable iSCSI, Fibre Channel SAN, or a file server based in Windows Server 2012 R2) and frees up the Hyper V host processors to carry more virtual machine workloads.

Without ODX, customers can saturate bandwidth on network connections, and utilize increased levels of CPU and Memory to perform certain data-related tasks, such as a large file copy, or a VM storage migration. These tasks can also take significant amounts of time, even on fast 10GbE networks meaning there may be performance degradation for a period of time, until the task is complete.

Offloaded data transfer (ODX) in Windows Server 2012 R2 enables you to accomplish more with your existing external storage arrays by letting you quickly move large files and virtual machines directly between storage arrays, which reduces host CPU and network resource consumption. When used with offload-capable SAN storage hardware, ODX lets a storage device perform a file copy operation without the main processor of the Hyper V host actually reading the content from one storage place and writing it to another.

ODX uses a token-based mechanism for reading and writing data within or between intelligent storage arrays. Instead of routing the data through the host, a small token is copied between the source and destination. The token simply serves as a point-in-time representation of the data.

As an example, when you copy a file or migrate a virtual machine between storage locations (either within or between storage arrays), a token that represents the virtual machine file is copied, which removes the need to copy the underlying data through the servers. In a token-based copy operation, the steps are as follows (see the following figure):




Figure 5 - Token-based copy operation

When a user attempts to copy a file from one volume to another that is residing on a SAN which supports **ODX**, the following happens automatically (even if copying and pasting through Explorer).

1. A user copies or moves a file by using Windows Explorer, a command line interface, or as part of a virtual machine migration.
2. Windows Server 2012 automatically translates this transfer request into an ODX (if supported by the storage array), and it receives a token that represents the data.
3. The token is copied between the source server and destination server.
4. The token is delivered to the storage array.
5. The storage array internally performs the copy or move and provides status information to the user.

### Requirements

ODX support in Hyper-V requires the following:

- Windows Server 2012, Windows Server 2012 R2, Hyper-V Server 2012 or Hyper-V Server 2012 R2
- ODX-capable hardware is required to host virtual hard disk files connected to the virtual machine as virtual SCSI devices or directly attached (sometimes referred to as pass-through disks).
- This optimization is also supported for natively attached, VHDX-based virtual disks.
- **VHD or VHDX-based virtual disks attached to virtual IDE do not support this optimization because integrated development environment (IDE) devices lack ODX support.**

### Why This Matters

**ODX frees up the main processor to handle virtual machine workloads, enabling native-like performance when your virtual machines read from and write to storage. When copying files, not only is the time taken to perform the task, significantly reduced, those copy operations don't consume valuable host resources, helping to ensure that a virtualized workload operates as efficiently as it would in a non-virtualized environment. From a virtual disk**

**perspective, the ODX capabilities ensure faster performance of crucial maintenance tasks for virtual hard drives (such as merge, move, and compact) that depend on copying large amounts of data without using processor time.**

## Enhanced Networking Performance

Windows Server 2012 provided a number of key networking capabilities specifically aimed at enhancing networking performance at both the host, and VM levels, many of which, integrated deeply with hardware innovation from our Partners, to drive this increased level of performance. With Windows Server 2012 R2, this has been enhanced even further, ensuring that Hyper-V is the optimal choice for virtualizing network-intensive workloads.

### Virtual Receive Side Scaling

Prior to 10 gigabit networking, one modern processor was usually more than enough to handle the networking workload of a VM. With the introduction of 10Gb/s NICs life became more complicated as the amount of data being sent to and received from a VM exceeded what a single processor could effectively handle. Our performance investigations found that since all network traffic was being processed on a single VP, a single VM was limited to (on average) 5Gbps, far below the full potential of the hardware installed in the system. The image below is a screen shot of the Task Manager from within a VM. In the figure below, VP3 is clearly being fully utilized and cannot support any additional traffic processing, even though it has 8 VP allocated.




Figure 6 – A single vCPU being fully utilized processing network traffic

Fortunately this problem was not new. Prior to this release, there was a similar situation with the introduction of multi-core machines for physical workloads. That experience had produced **Receive Side Scaling (RSS)**. RSS spreads traffic from the network interface card (NIC), based on TCP flows, and to multiple processors for simultaneous processing of TCP flows. I will not go in to the details of RSS but inquiring minds can read more on RSS at this link, [Receive Side Scaling](#). This enabled physical workloads to optimally utilize bandwidth and cores available.

Similar to how RSS distributes networking traffic to multiple cores in physical machines, **vRSS** spreads networking traffic to multiple VPs in each VM by enabling RSS inside the VM. With vRSS enabled, a VM is able to process traffic on multiple VPs simultaneously and increase the amount of throughput it is able to handle.

vRSS is managed in the VM the same way RSS is managed on a physical machine. In the VM open a PowerShell instance with administrator rights. Type the following cmdlet and substitute your network connection in the -Name field or simply use "\*" to enable across all adapters.

```
PS C:\> Enable-NetAdapterRss -Name "Ethernet"
```

Rerunning the same test as shown above now gives much improved results. Again, the Task Manager from the VM is shown. The processing is now distributed to all the VPs and the VM is handling 9.8 Gbps of network traffic, double the previous throughput and effectively line rate on the 10G NIC. The best part about this new feature is it doesn't require anyone to install or replace any hardware; this was all done by maximizing the use of existing resources in the server.




Figure 7 – All vCPUs being utilized to process network traffic

One thing to be aware of, is that vRSS is not enabled by default on any VMs. There are extra calculations that must be done to accomplish the spreading which leads to higher CPU utilization in the host. This means that small VMs with minimal or average network traffic will not want to enable this feature. This feature is meant for VMs that process high network traffic like file servers or gateways.

## Requirements

To utilize vRSS inside a virtual machine, you'll need the following:

- Windows Server 2012 R2 Hyper-V or Hyper-V Server 2012 R2
- VMQ-capable NIC for the Hyper-V Switch
- RSS enabled within the Guest OS NIC Properties

## Why This Matters

**In the past, VMs might have trouble achieving network throughput approaching 10Gbps due to the processing load on a single CPU core. vRSS alleviates this problem by spreading the processing across multiple cores on the host and multiple cores on the VM. This allows VMs to sustain a greater networking traffic load and gives customer the confidence they need to virtualize those network-intensive workloads.**

## Dynamic Virtual Machine Queue

The **Virtual Machine Queue (VMQ)** is a hardware virtualization technology for the efficient transfer of network traffic to a virtualized host OS. A VMQ capable NIC classifies incoming frames to be routed to a receive queue based on filters which associate the queue with a VM's virtual NIC. Each virtual machine device buffer is assigned a VMQ, which avoids needless packet copies and route lookups in the virtual switch.

Essentially, VMQ allows the host's single network adapter to appear as multiple network adapters to the virtual machines, allowing each virtual machine its own dedicated network adapter. The result is less data in the host's buffers and an overall performance improvement to I/O operations.

These hardware queues may be affinity-tuned to different CPUs thus allowing for receive scaling on a per-VM NIC basis. Windows Server 2008 R2 allowed administrators to statically configure the number of processors available to process interrupts for VMQ.




Figure 8 – Hyper-V Host Processing Network Traffic on CPU0 with no VMQ Enabled

**Without VMQ** - the Hyper-V virtual switch is responsible for routing and sorting of packets that are inbound to the VMs. This can lead to a lot of CPU processing for the virtual switch on heavily loaded Hyper-V hosts. Without VMQ technology and RSS, a majority of the network processing would burden CPU0 and would ultimately limit the scale of the solution.




Figure 9 – Hyper-V Host Processing Network Traffic across multiple cores with VMQ Enabled

**With VMQ** - When VMQ is enabled, a dedicated queue is established on the physical network adapter for each virtual network adapter that has requested a queue. As packets arrive for a virtual network adapter, the physical network adapter places them in that network adapter's queue. When packets are indicated up, all the packet data in the queue is delivered directly to the virtual network adapter. Packets arriving for virtual network adapters that don't have a dedicated queue, as well as all multicast and broadcast packets,

are delivered to the virtual network in the default queue. The virtual network handles routing of these packets to the appropriate virtual network adapters as it normally would. This reduces a significant amount of CPU overhead on the host associated with network traffic as it spreads the load over multiple cores.




Figure 10 – Hyper-V Host Processing Network Traffic Dynamically Across Multiple Cores with DVMQ

**With Dynamic VMQ** – First introduced in Windows Server 2012, DVMQ dynamically distributes incoming network traffic processing to host processors (based on processor usage and network load). In times of heavy network load, Dynamic VMQ automatically recruits more processors. In times of light network load, Dynamic VMQ relinquishes those same processors. VMQ spreads interrupts for network traffic across available processors. In Windows Server 2012 and Windows Server 2012 R2, the Dynamic VMQ capability allows an adaptive algorithm to modify the CPU affinity of queues without the need of removing/re-creating queues. This results in a better match of network load to processor use, resulting in increased network performance.

### Requirements

To utilize Dynamic VMQ, you'll need the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2
- VMQ-capable NIC for the Hyper-V Switch

### Why This Matters

**To drive the highest levels of performance for virtualized workloads, the underlying host OS needs to operate as efficiently as possible, especially when under pressure from multiple heavily-laden virtual machines. With Dynamic VMQ, customers can take advantage of VMQ capable network cards to more efficiently process the network traffic utilizing the different cores in the host system. With this being dynamic in Windows Server 2012 R2, this ensures the best use of resources on the host, and ensures the virtual machines operate as efficiently as possible.**

### Single-Root I/O Virtualization

Single Root I/O Virtualization (**SR-IOV**) is an open standard introduced by the PCI-SIG, the special-interest group that owns and manages PCI specifications. SR-IOV works in conjunction with system chipset support



for virtualization technologies that provide remapping of interrupts and Direct Memory Access, and enables assignment of SR-IOV-capable devices directly to a virtual machine.

As discussed earlier, customers are looking to virtualize their most network-intensive workloads, yet without capabilities like **vRSS** and **Dynamic VMQ**, it can provide very difficult to achieve the same kind of performance you'd expect in the physical world, inside VMs. This is typically due to the CPU overhead of processing NIC interrupts and DMA requests, along with heavy Hyper-V Switch activity which also contributes towards higher levels of CPU usage.

Hyper-V in Windows Server 2012 introduced support for SR-IOV-capable network devices and lets an SR-IOV virtual function of a physical network adapter be assigned directly to a virtual machine. This increases network throughput and reduces network latency while also reducing the host CPU overhead that is required for processing network traffic. The following figure shows the architecture of SR-IOV support in Hyper-V.

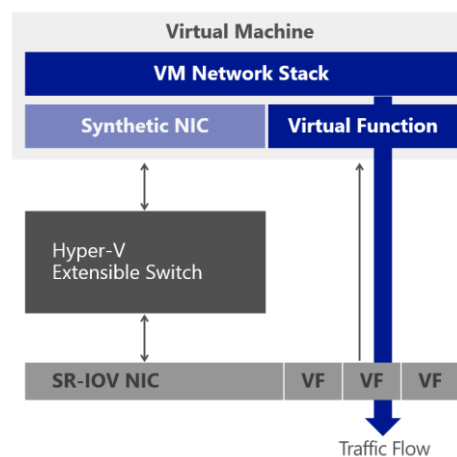


Figure 11 – SR-IOV High Level Architecture – Virtual Function enabled within the VM

Essentially, SR-IOV works by remapping bypassing the Hyper-V Extensible Switch and mapping virtual network functions from the physical NIC directly to the VM, in essence, bypassing the Hyper-V Switch.

### SR-IOV & Live Migration

In this scenario, we will walk through how the SR-IOV configuration is handled during a live migration of a VM from one host to another. After the Hyper-V guest is started, network traffic flows over the synthetic data path. **If the physical network adapter supports the single root I/O virtualization (SR-IOV) interface, it can enable one or more PCI Express (PCIe) Virtual Functions (VFs).** Each VF can be attached to a Hyper-V child partition. When this happens, network traffic flows over the hardware-optimized SR-IOV VF Data Path.

When a Live Migration is initiated, a VF failover occurs, from SR-IOV VF, to the native synthetic data path, and traffic then flows over the synthetic data path, out through the Hyper-V Switch.

The transition between the VF and synthetic data paths occurs with minimum loss of packets and prevents the loss of TCP connections.

The VM is then live migrated, during which time, the VM remains accessible on the network, but always via the synthetic data path. Upon reaching the new target host, should the new host have SR-IOV functionality enabled, as per the source host, the VM will automatically be assigned a VF from the SR-IOV NIC, and traffic will failover to pass via the VF inside the VM, as oppose to the synthetic data path. If however, the new host does not have SR-IOV capable hardware, the VM will simply continue to utilize the synthetic data path.

### Requirements

To utilize SR-IOV within a virtual machine, you'll need the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2
- A physical host system that supports SR-IOV (such as Intel VT-d2), including chipset support for interrupt and DMA remapping and proper firmware support to enable and describe the platform's SR-IOV capabilities to the operating system. SR-IOV may need to be enabled in the BIOS.
- An SR-IOV-capable network adapter and driver in both the management operating system (which runs the Hyper-V role) and each virtual machine where a virtual function is assigned.

### Why This Matters

**SR-IOV is another example of deep integration with hardware innovation. By integrating with hardware investments, customers can protect their investment, whilst driving even higher levels of performance. In this case, networking performance, resulting in higher bandwidth and throughput to a VM, and reduced latency and CPU usage.**

**It's important for customers also, that with Hyper-V, we provide the highest levels of performance, but don't sacrifice agility. This is why we ensured that SR-IOV worked seamlessly with Live Migration, ensuring that even network-intensive workloads can be migrated without loss of TCP connectivity.**

## Enhanced Resource Management

Windows Server 2012 provided administrators with a number of new and powerful capabilities to manage and allocate resources quickly and efficiently. From controlling virtual machine memory to increase host density, through to ensuring SLAs can be met with granular bandwidth management controls, these powerful capabilities helped IT administrators optimize their infrastructure. In Windows Server 2012 R2, these capabilities have been refined, and enhanced to provide even greater levels of capability for the IT administrator.

### Dynamic Memory

Dynamic Memory, introduced in Windows Server 2008 R2 with SP1, helps IT administrators manage physical memory more efficiently. With Dynamic Memory, Hyper-V treats memory as a shared resource that can be automatically reallocated among running virtual machines. Dynamic Memory adjusts the amount of memory available to a virtual machine based on changes in memory demand and values that the IT administrator specifies.

Dynamic Memory in Windows Server 2008 R2 Hyper-V, included two key settings. The first was known as "startup memory," which was defined as the minimum amount of memory that a virtual machine could have. Windows however, typically requires more memory during startup than at steady state. The second setting was "maximum memory", which, as the name suggests, was the maximum amount of memory that the VM would ever receive at any one time. These were both static values, in the sense that you could set them whilst the VM was switched off, but as soon as the VM was powered on, the settings remained as they were.

This locking of the settings meant that a memory upgrade would require shutting down the virtual machine. This is a common challenge for administrators who frequently have to upgrade the maximum

amount of memory for a virtual machine as demand increases. For example, consider a virtual machine running SQL Server and configured with a maximum of 8 GB of RAM. With an increase in the size of the databases and increased transactions, the virtual machine now requires more memory. In Windows Server 2008 R2 with SP1, you must shut down the virtual machine to perform the upgrade, which requires planning for downtime and decreasing business productivity.

Fast-growing organizations whose workloads are rapidly expanding often need to add more virtual machines to their hosts. These organizations want to optimize the number of virtual machines they can place on a host server to minimize the number of expensive host servers that they need. With the Hyper-V Dynamic Memory improvements in Windows Server 2012 R2, IT administrators can now allocate virtual machine memory resources more efficiently and dramatically increase virtual machine consolidation ratios.

With Windows Server 2012 R2, Dynamic Memory has a new configuration option, "minimum memory." **Minimum memory** lets the administrator specify a value lower than the "startup memory", thus allowing Hyper-V to reclaim the unused memory from the virtual machines after startup. This can result in increased virtual machine consolidation numbers, especially in VDI environments.

In addition, these settings are no longer locked, meaning an administrator can adjust, whilst the VM is running, both the **minimum** and **maximum VM memory**. This means that in the example we discussed earlier, with a database with growing demand, the IT Administrator can increase the maximum memory for that particular VM, allowing it to meet the increased demand, all without downtime.

### Smart Paging

Windows Server 2012 also introduces **Hyper-V Smart Paging** for robust virtual machine restart. Although minimum memory increases virtual machine consolidation numbers, it also brings a challenge. If a virtual machine has a smaller amount of memory than its startup memory and it is restarted, Hyper-V needs additional memory to restart the machine. Due to host memory pressure or virtual machines' states, Hyper-V may not always have additional memory available. This can cause sporadic virtual machine restart failures in customer environments. In Windows Server 2012 R2, Hyper-V Smart Paging is used to bridge the memory gap between minimum memory and startup memory and let virtual machines restart reliably.

**Hyper-V Smart Paging** is a memory management technique that uses disk resources as additional, temporary memory when more memory is required to restart a virtual machine. This approach has both advantages and drawbacks. It provides a reliable way to keep the virtual machines running when no physical memory is available. However, it can degrade virtual machine performance because disk access speeds are much slower than memory access speeds.

To minimize the performance impact of Smart Paging, Hyper-V uses it only when all of the following occur:

- The virtual machine is being restarted.
- No physical memory is available.
- No memory can be reclaimed from other virtual machines that are running on the host.

Hyper-V Smart Paging is not used when:

- A virtual machine is being started from an off state (instead of a restart).
- Oversubscribing memory for a running virtual machine would result.
- A virtual machine is failing over in Hyper-V clusters.

Hyper-V continues to rely on internal guest paging when host memory is oversubscribed because it is more effective than Hyper-V Smart Paging. With internal guest paging, the paging operation inside virtual machines is performed by Windows Memory Manager. Windows Memory Manager has more information than the Hyper-V host, about memory use within the Guest OS, which means it can provide Hyper-V with better information to use when it chooses the memory to be paged. Because of this, internal guest paging incurs less overhead to the system than Hyper-V Smart Paging.

The figure on the following page shows the mapping of memory for a virtual machine that is being restarted by using Hyper-V Smart Paging.

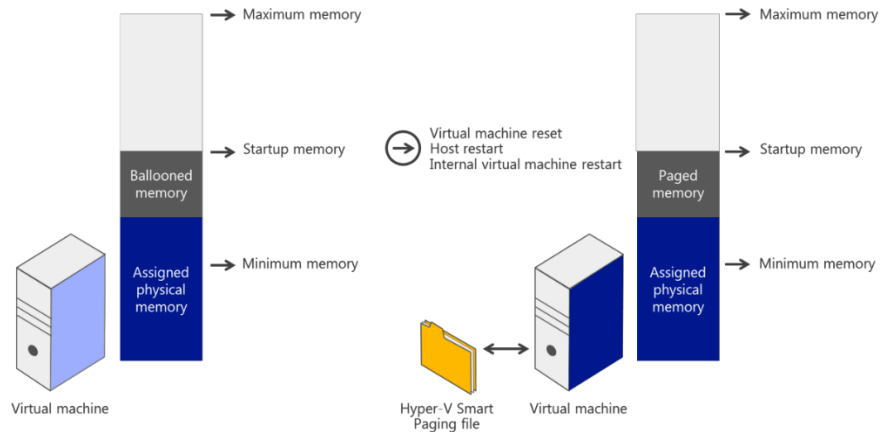


Figure 12 – Hyper-V Smart Paging

To further reduce the impact of Hyper-V Smart Paging, after a virtual machine completes the startup process, Hyper-V removes memory from the virtual machine, coordinating with Dynamic Memory components inside the guest (a process sometimes referred to as “ballooning”), so that the virtual machine stops using Hyper-V Smart Paging. With this technique, use of Hyper-V Smart Paging is temporary, and is not expected to be longer than 10 minutes.

The following figure shows Hyper-V removing memory from the virtual machine after it completes the startup process.

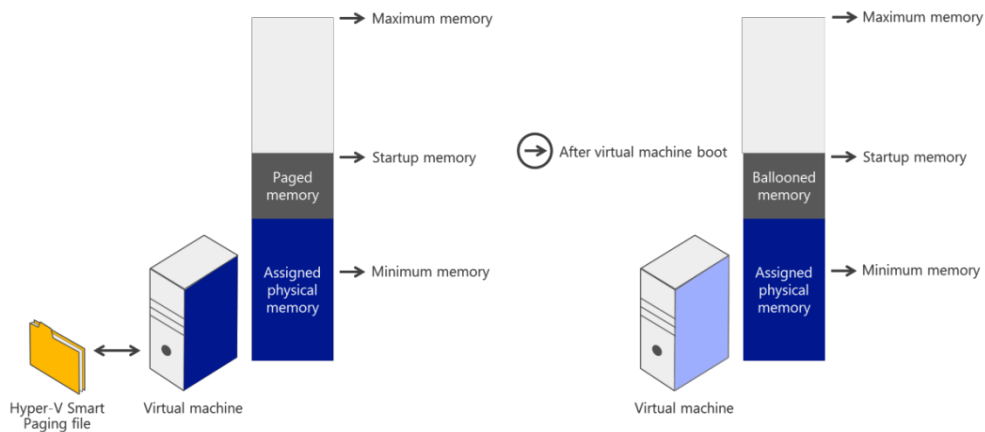


Figure 13 – Removing paged memory after virtual machine restart

Also note the following about how Hyper-V Smart Paging is used:

- Hyper-V Smart Paging files are created only when needed for a virtual machine.
- After the additional amount of memory is removed, Hyper-V Smart Paging files are deleted.

- Hyper-V Smart Paging is not used for this virtual machine again until another restart occurs and not enough physical memory exists.

## Requirements

Dynamic Memory requires the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2

### Why This Matters

**Dynamic Memory in Windows Server 2012 R2 Hyper-V help you reach higher consolidation numbers with improved reliability of Hyper-V operations. You can make memory configuration changes for your virtual machines without shutting down the virtual machines. If you have idle or low-load virtual machines, as in pooled VDI environments, Dynamic Memory additions in Hyper-V let you increase consolidation and improve reliability for restart operations. This can lead to lower costs for customers. With runtime configuration changes for Dynamic Memory, overall IT productivity is expected to increase with reduced downtime and increased agility to respond to requirement changes. You also gain agility in responding to requirement changes with these new capabilities**

## Hyper-V Resource Metering

Your computing resources are limited. You need to know how different workloads draw upon these resources, even when they are virtualized. In Windows Server 2012 R2 Hyper-V, Resource Metering, helps you track historical data on the use of virtual machines and gain insight into the resource use of specific servers. You can use this data to perform capacity planning, to monitor consumption by different business units or customers, or to capture data needed to help redistribute the costs of running a workload. You could also use the information that this feature provides to help build a billing solution, so that customers of your hosting services can be charged appropriately for resource usage.

### Metrics

Windows Server 2012 R2 offers two ways to obtain historical data on customer usage of virtual machine resources: Hyper-V cmdlets in Windows PowerShell, and the new APIs in the Virtualization WMI Provider.

Hyper-V exposes the metrics in the following table for resource use:

Metric	Units	Description
Average CPU Use	Megahertz (MHz)	The average amount of CPU used by a virtual machine over a period of time
Average Memory Use	Megabytes (MB)	The average amount of physical memory used by a virtual machine over a period of time
Minimum Memory Use	Megabytes (MB)	The lowest amount of physical memory assigned to a virtual machine over a period of time
Maximum Memory Use	Megabytes (MB)	The highest amount of physical memory assigned to a virtual machine over a period of time

Metric	Units	Description
Maximum Disk Allocation	Megabytes (MB)	The highest amount of disk space capacity allocated to a virtual machine over a period of time
Incoming Network Traffic	Megabytes (MB)	Total incoming network traffic, for a virtual network adapter over a period of time
Outgoing Network Traffic	Megabytes (MB)	Total outgoing network traffic for a virtual network adapter over a period of time
Total Average Normalized IOPS	IOPS	Average throughput of normalized I/O operations per second over a period of time
Total Average Latency	Milliseconds (MS)	Average storage latency for a virtual machine recorded over a period of time
Total Data Written	Megabytes (MB)	Total data written to disk for a virtual machine recorded over a period of time
Total Data Read	Megabytes (MB)	Total data read from disk for a virtual machine recorded over a period of time

Table 2 – Hyper-V Resource Metering Metrics

### Use of Network Metering Port ACLs

Enterprises pay for the Internet traffic coming into and going out of their datacenters, but not for the network traffic within the datacenters. For this reason, providers generally consider Internet and intranet traffic separately for the purposes of billing. To differentiate between Internet and intranet traffic, providers can measure incoming and outgoing network traffic for any IP address range, by using Network Metering Port ACLs.

### Metering VM Usage in a Multitenant Environment

Hyper-V in Windows Server 2012 R2 lets providers build a multitenant environment in which virtual machines can be served to multiple clients in a more isolated and secure way, as shown in the following figure. Because a single client may have many virtual machines, aggregation of resource use data can be challenging. However, Windows Server 2012 R2 simplifies this task by using resource pools, a feature available in Hyper-V. Resource pools are logical containers that collect resources of virtual machines belonging to one client, permitting single-point querying of the client's overall resource use.

The following figure is an example of Resource Metering in a two-tenant environment that is built with Hyper-V in Windows Server 2012 R2.




Figure 14 – Two-tenant environment built with Hyper-V in Windows Server 2012 R2

The following figure shows a basic model of Resource Metering.

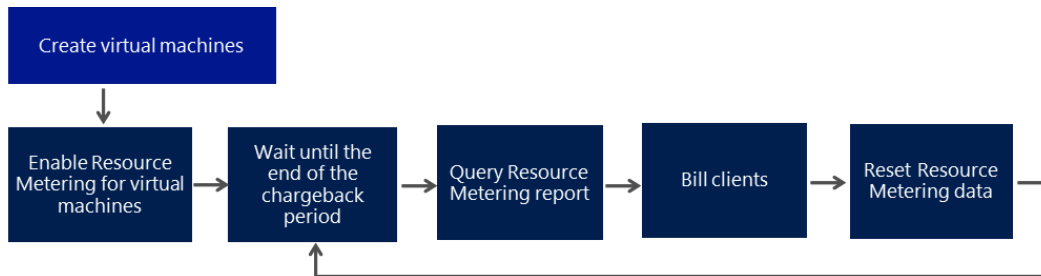


Figure 13 – Resource Metering Model

In this model, a hosting provider does the following:

1. Create virtual machines for a customer and enable Resource Metering once for the virtual machines.
2. In a multitenant environment, the provider would enable metering on each resource pool. Hyper-V then tracks resource use for each virtual machine until that virtual machine is deleted.
3. Query resource use data at the end of each chargeback period, and use this data to bill clients as needed.
4. Reset the data at the end of each chargeback period, so that Hyper-V can begin tracking resource use for the new chargeback period.

Resource Metering works with all Hyper-V operations. Movement of virtual machines between Hyper-V hosts (such as through live, offline, or storage migration) does not affect the collected data.

### Why This Matters

The Resource Metering feature in Windows Server 2012 R2 Hyper-V makes it easier for you to track historical data about each customer’s use of virtual machines. Through resource pools, which are part of this technology, Hyper-V lets providers’ aggregate usage data in a multitenant environment, in which each customer or business unit may have many virtual machines. With this feature, you can perform capacity planning or monitor resource

**consumption by various business units or customers. Third-party ISVs can use data provided by this feature to build more reliable, cost-effective, usage-based billing solutions.**

## Network Quality of Service (QoS)

Public cloud hosting providers and large enterprises must often run multiple application servers on servers running Hyper-V. Hosting providers that host customers on a server running Hyper-V must deliver performance that's based on service level agreements (SLAs). Enterprises want to run multiple application servers on a server running Hyper-V with the confidence that each one will perform predictably.

Most hosting providers and enterprises use a dedicated network adapter and a dedicated network for a specific type of workload, such as storage or live migration, to help achieve network performance isolation on a server running Hyper-V. This strategy works for 1-gigabit Ethernet (GbE) network adapters, but becomes impractical for those using or planning to use 10 GigE network adapters.

For most deployments, one or two 10 GigE network adapters provide enough bandwidth for all the workloads on a server running Hyper-V. However, 10-GbE network adapters and switches are considerably more expensive than their 1-GbE counterparts. To optimize the 10 GigE hardware, a server running Hyper-V requires new capabilities to manage bandwidth.

Windows Server 2012 R2 expands the power of QoS by providing the ability to assign a minimum bandwidth to a virtual machine or a service. This feature is important for hosting companies who need to honor SLA clauses that promise a minimum network bandwidth to customers. It's equally important to enterprises that require predictable network performance when they run virtualized server workloads on shared hardware.

In Windows Server 2008 R2, QoS supports the enforcement of maximum bandwidth. This is known as rate limiting. Consider a typical physical server running Hyper-V in which the following four types of network traffic share a single 10 GigE network adapter:

1. Traffic between virtual machines and resources on other servers.
2. Traffic to and from storage.
3. Traffic for live migration of virtual machines between servers running Hyper-V.
4. Traffic to and from a CSV (intercommunication between nodes in a cluster).

If virtual machine data is rate-limited to 3 gigabits per second (Gbps), the sum of the virtual machine data throughputs cannot exceed 3 Gbps at any time, even if other network traffic types don't use the remaining 7 Gbps of bandwidth. However, this also means that the other types of traffic can reduce the actual amount of bandwidth available for virtual machine data to unacceptable levels, depending on how their maximum bandwidths are defined.

### Minimum Bandwidth

Windows Server 2012 R2 provides a minimum bandwidth setting. Unlike maximum bandwidth, which is a bandwidth cap, minimum bandwidth is a bandwidth floor. It assigns a certain amount of bandwidth to a given type of traffic. The following figure shows how minimum bandwidth works for each of the four types of network traffic flows in three different time periods: T1, T2, and T3.



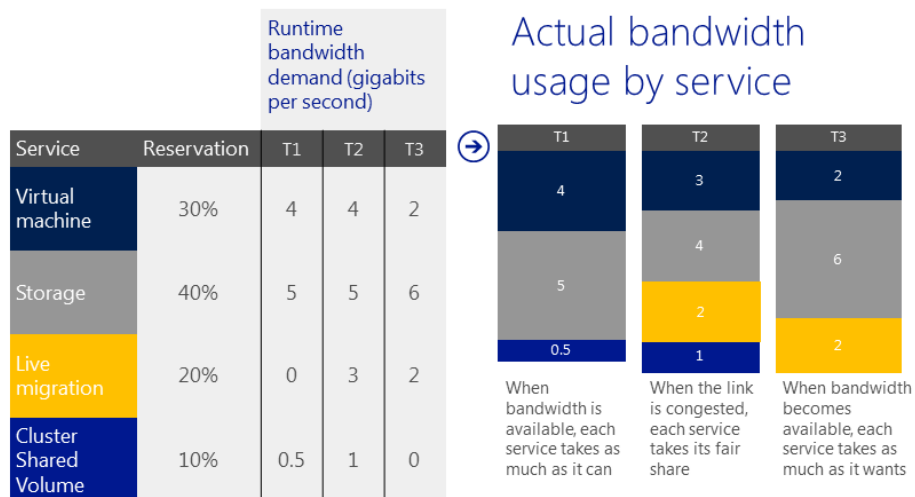


Figure 15 – Assigning minimum bandwidth to services

In this figure, the table on the left shows the configuration of the minimum amount of required bandwidth that a given type of network traffic flow needs. For example, storage is configured to have at least 40 percent of the bandwidth (4 Gbps of a 10 GigE network adapter) at any time. The table on the right shows the actual amount of bandwidth that each type of network traffic has in T1, T2, and T3. In this example, storage is actually sent at 5 Gbps, 4 Gbps, and 6 Gbps, respectively, in the three periods.

The characteristics of minimum bandwidth can be summarized as follows:

- In the event of congestion, when the desired network bandwidth exceeds the available bandwidth (such as in the T2 period in the figure), minimum bandwidth helps ensure that each type of network traffic receives up to its assigned bandwidth. For this reason, minimum bandwidth is also known as fair sharing. This characteristic is essential for converging multiple types of network traffic on a single network adapter.
- If there's no congestion—that is, when sufficient bandwidth is available to accommodate all network traffic (such as in the T1 and T3 periods)—each type of network traffic can exceed its quota and consume as much bandwidth as is available. This characteristic makes minimum bandwidth superior to maximum bandwidth in using available bandwidth.

### Relative Minimum Bandwidth

If the importance of workloads in virtual machines is relative, you can use relative minimum bandwidth, where you assign a weight to each virtual machine, giving the more important ones a higher weight. You determine the bandwidth fraction that you assign to a virtual machine by dividing the virtual machine's weight by the sum of all the weights of virtual machines that are attached to the Hyper-V Extensible Switch. The following figure illustrates relative minimum bandwidth

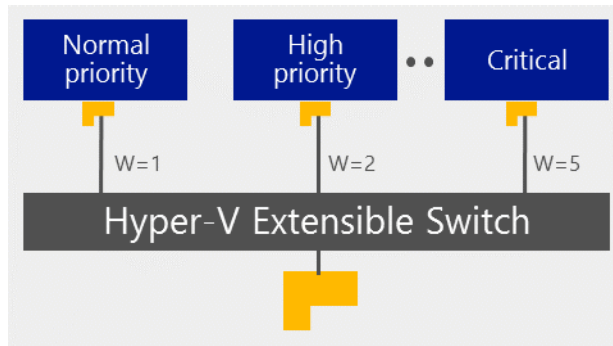


Figure 16 – Relative minimum bandwidth

### Strict Minimum Bandwidth

If you want to provide an exact bandwidth, you should use strict minimum bandwidth where you assign an exact bandwidth quota to each virtual machine that is attached to the Hyper-V Extensible Switch.




Figure 17 – Strict minimum bandwidth

### Minimum Bandwidth Mechanisms

Windows Server 2012 R2 offers two different mechanisms to enforce minimum bandwidth: one in software (through the newly enhanced packet scheduler in Windows), and the other through the network adapters that support Datacenter Bridging (DCB). In both cases, network traffic must be classified first. Windows either classifies a packet itself or gives instructions to a network adapter to classify it. The result of classification is a number of traffic flows in Windows, and a given packet can belong to only one of them.

For example, a traffic flow could be a live migration connection, a file transfer between a server and a client, or a Remote Desktop connection. Based on how the bandwidth policies are configured, either the packet scheduler in Windows or the network adapter dispatches the packets at a rate equal to, or higher than, the minimum bandwidth configured for the traffic flow.

Each of the two mechanisms has its own advantages and disadvantages:

- The software solution, which is built on the new packet scheduler in Windows Server 2012 R2, provides a fine granularity of classification. It is the only viable choice if many traffic flows require minimum bandwidth enforcement. A typical example is a server running Hyper-V hosting many virtual machines, where each virtual machine is classified as a traffic flow.
- The hardware solution, which depends on DCB support on the network adapter, supports far fewer traffic flows. However, it can classify network traffic that doesn't originate from the networking stack. A typical scenario involves a CNA that supports iSCSI offload, in which iSCSI traffic bypasses the networking stack and is framed and transmitted directly by the CNA. Because the packet scheduler in the networking stack doesn't process this offloaded traffic, DCB is the only viable choice to enforce minimum bandwidth.

You can employ both of these mechanisms on the same server. For example, a server running Hyper-V has two physical network adapters: one that binds to a virtual switch and serves virtual machine data, and another that serves the rest of the traffic of the host server. You can enable the software-based minimum bandwidth in Hyper-V to help ensure bandwidth fair sharing among virtual machines, and enable the hardware-based minimum bandwidth on the second network adapter to help ensure bandwidth fair sharing among various types of network traffic from the host server.

It is not recommended that you enable both mechanisms at the same time for a given type of network traffic, however. For example, consider live migration and storage traffic that are configured to use the second network adapter on the server running Hyper-V. If you have already configured the network adapter to allocate bandwidth for live migration and storage traffic using DCB, you shouldn't also configure the packet scheduler in Windows to do the same, and vice versa.

### **Configuring and Managing QoS**

In Windows Server 2012 R2, you manage QoS policies and settings dynamically with Windows PowerShell. The new QoS cmdlets support both the QoS functionalities that are available in Windows Server 2008 R2—such as maximum bandwidth and priority tagging—and the new features such as minimum bandwidth that are available in Windows Server 2012 and Windows Server 2012 R2.

#### **Requirements**

Minimum QoS can be enforced through the following two methods:

- The first method relies on software built into Windows Server 2012 and has no other requirements.
- The second method, which is hardware-assisted, requires a network adapter that supports DCB.
- Both options require Windows Server 2012, Windows Server 2012 R2, Hyper-V Server 2012, or Hyper-V Server 2012 R2.

For hardware-enforced minimum bandwidth, you must use a network adapter that supports DCB and the miniport driver of the network adapter must implement the NDIS QoS APIs. A network adapter must support Enhanced Transmission Selection (ETS) and Priority-Based Flow Control (PFC) to pass the NDIS QoS logo test created for Windows Server 2012. Explicit Congestion Notification (ECN) is not required for the logo. The IEEE Enhanced Transmission Selection (ETS) specification includes a software protocol called Data Center Bridging Exchange (DCBX) to let a network adapter and switch exchange DCB configurations. DCBX is also not required for the logo.

Enabling QoS in Windows Server 2012 R2, when it is running as a virtual machine, is not recommended. Minimum bandwidth enforced by the packet scheduler works optimally on 1 GbE or 10 GigE network adapters.

#### **Why This Matters**

**To help improve performance in virtualized environments, Windows Server 2012 R2 provides powerful QoS bandwidth management features which enable you to assign a minimum bandwidth to a virtual machine or a service. Hosting providers and enterprises can now optimize the number of virtual machines on their Hyper-V servers and have confidence that they will perform as expected. This helps to ensure that customers won't be affected or compromised by other customers on their shared infrastructure, which includes computing,**

**storage, and network resources. Hosters and Enterprises can also sandbox their applications and provide different SLAs/pricing depending on bandwidth guarantees.**

## Storage Quality of Service (QoS)

Starting in Windows Server 2012 R2, Hyper-V includes the ability to set certain quality-of-service (QoS) parameters for storage on the virtual machines. Storage QoS provides storage performance isolation in a multitenant environment and mechanisms to notify you when the storage I/O performance does not meet the defined threshold to efficiently run your virtual machine workloads.

Storage QoS provides the ability to specify a maximum input/output operations per second (IOPS) value for an individual virtual hard disk. Similar to how Network QoS can be used to limit a noisy VM in terms of network traffic and utilization, an administrator can throttle the storage I/O to stop a tenant from consuming excessive storage resources that may impact another tenant. Storage QoS supports Fixed, Dynamic and Differencing virtual hard disks.

Imagine a scenario where a virtual machine is running an intensive database workload, which is stored on a secondary virtual hard disk.




Figure 18 – VM with intensive workload using significant IOPS on Data VHDX

With Storage QoS enabled, the administrator can cap the incoming and outgoing IOPS to that virtual disk, alleviating the pressure on the SAN, and freeing up resources that can be used by other virtual machines on that host.




Figure 19 – VM with intensive workload, with IOPS capped at 500 with Storage QoS

An administrator can also set a minimum IOPS value. They will be notified when the IOPS to a specified virtual hard disk is below a threshold that is needed for its optimal performance.

The virtual machine metrics infrastructure is also updated, with storage related parameters to allow the administrator to monitor the performance and chargeback related parameters. Maximum and minimum values are specified in terms of normalized IOPS where every 8 K of data is counted as an I/O.

With these key capabilities, the administrator can more easily plan for, and gain acceptable performance from the investment in storage technologies.

### Requirements

For Storage QoS, you will require the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012 or Hyper-V Server 2012 R2

### Why This Matters

**Similar to Network QoS, To help improve overall performance in virtualized environments, Windows Server 2012 R2 Hyper-V, with Storage QoS, helps to control the balance of VM storage demand, and storage performance capacity. No SAN has unlimited resource in terms of IOPS, and with Storage QoS, administrator can manage the allocation of bandwidth down to the granularity of a virtual disk. This helps to prevent noisy VMs consuming more than a fair share of storage resources, and helps to prevent contention within the infrastructure.**

## Virtualized Workload Performance

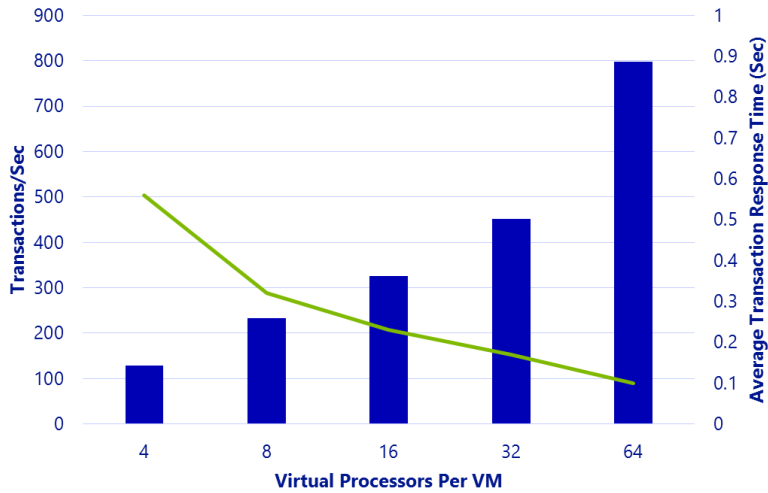
Windows Server 2012 R2 Hyper-V can scale to meet the demands of your most intensive workloads. We've looked at a number of the core features of Hyper-V that unlock the highest levels of performance. From architectural capabilities such as NUMA, to deep integration with hardware capabilities for powerful offloading, and from features such as **Dynamic Memory** and **Smart Paging** through to **Network** and **Storage QoS**. All of these capabilities help to ensure that when virtualizing your key workloads, they run at their best on Hyper-V.

Over the last 12 months, Microsoft has worked closely with Enterprise Strategy Group, who performed lab testing and analysis on a number of Microsoft key workloads, running virtualized on Windows Server 2012. These included SQL Server 2012, Exchange 2013 and SharePoint 2013.

### SQL Server 2012

Firstly, ESG tested an existing SQL Server 2012 OLTP workload that was previously vCPU limited. This test was performed previously on Windows Server 2008 R2 Hyper-V, which was restricted in terms of scale by the 4 vCPU per VM limit. With Windows Server 2012, and subsequently, R2, this limit has grown to 64 vCPUs per VM, as demonstrated in the below figure.

## Hyper-V Virtual CPU Scalability with OLTP Workloads

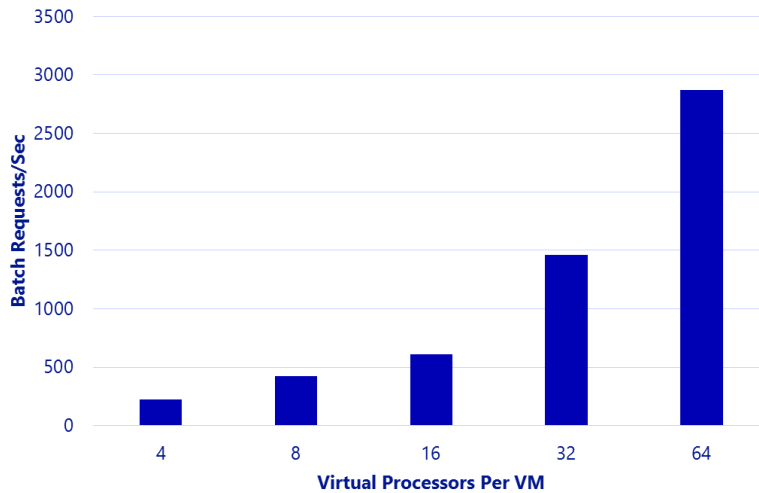


Windows Server 2012, SQL Server 2012, Single VM, 64GB of RAM

Figure 20 – Graph of a Hyper-V Virtual CPU Scalability with OLTP Workloads

With Hyper-V’s support for 64 vCPUs per VM, testing showed a 6x performance increase, with a 5x improvement in transaction response time over previous versions of Hyper-V. Additionally, ESG recorded the number of SQL Server Batch Requests, per second, that the Hyper-V VM could handle, with the results shown in the following graph:

## Hyper-V Enabled SQL Batch Request Scalability



Windows Server 2012, SQL Server 2012, Single VM, 64GB of RAM

Figure 21 – Graph of a Hyper-V Enabled SQL Batch Request Scalability

2,870 SQL Server batch requests per second were recorded during the 64 vCPU test. To put this into perspective, Microsoft documentation indicates that “over 1,000 batch requests per second indicate a very busy SQL Server”. Finally, ESG tested the performance of a physical SQL Server with a similarly configured virtual equivalent, as shown in the following graph:




Figure 22 – Graph of a Hyper-V OLTP Workload Analysis

The aim of this test was to quantify the manageably low difference in performance between the brokerage application running in a Hyper-V virtual machine and a native physical server. An OLTP workload running on a 75,000 brokerage customer database deployed in a Hyper-V virtual machine processed just over 6% fewer transactions per second compared to the same workload running on a similarly configured physical server.

### Exchange 2013

With Windows Server 2012 Hyper-V, ESG Lab performed hands-on testing of a virtualized tier-1 Exchange 2013 application workload. The workload used was designed to simulate thousands of Exchange users performing typical activities including sending and receiving e-mails, making calendar entries, updating contacts, and managing to-do lists. The graph below showcases the results from the testing:

## Exchange Workload Scalability on Windows Server 2012 with Hyper-V

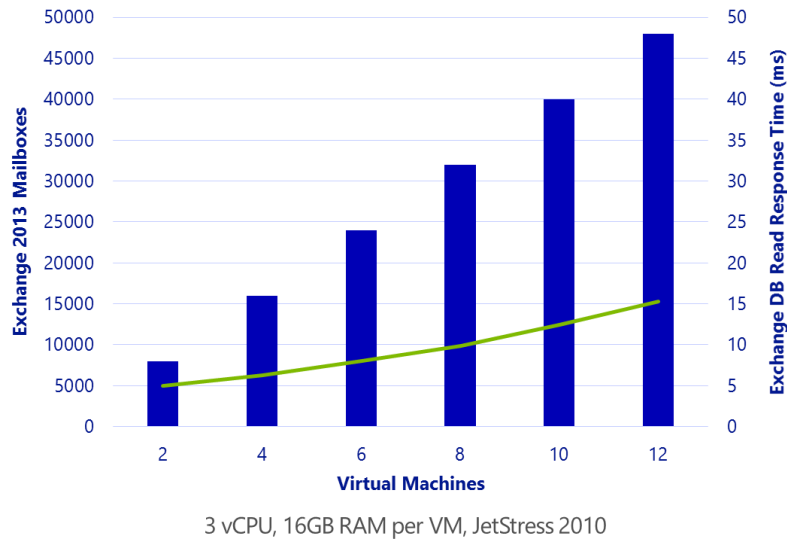


Figure 23 – Graph of virtualized Exchange 2013 scalability

As you can see from the results, an Exchange 2013 infrastructure deployed within 12 Hyper-V VMs, running on a single physical server, supported the I/O requirements of up to 48,000 simulated users, while average database read response times ranged between 5.02 and 15.31ms, well below the Microsoft recommended limit of 20 milliseconds.

### SharePoint 2013

With Windows Server 2012 Hyper-V, ESG Labs tested a virtualized configuration of SharePoint 2013 and their findings included that the performance, scalability, and low overhead of Hyper-V can be used to reduce costs while improving the manageability, flexibility, and availability of consolidated SharePoint 2013 workloads. The graph below showcases the results from the testing:

## SharePoint Workload Scalability on Windows Server 2012 with Hyper-V

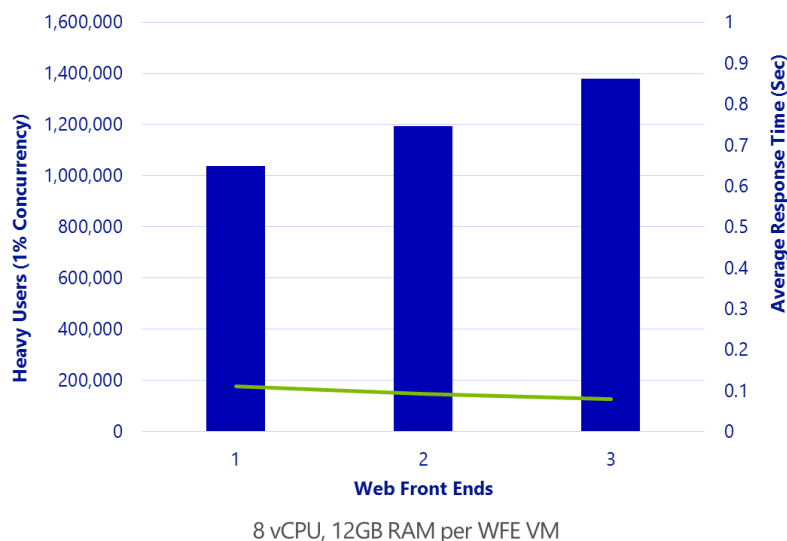




Figure 24 – Graph of virtualized SharePoint 2013 scalability

A SharePoint 2013 infrastructure deployed within 5 Hyper-V VMs (3 WFE, 1 App, 1 SQL), running on a single physical server, backed by SSD-based, mirrored Storage Spaces, supported the demand of over 1.3 million heavy users (60 requests per hour), with 1% concurrency, running a lightweight, non-blocking workload), with measurably low response times throughout.

The same configuration was retested, pushing the CPU utilization even higher, to see if higher numbers could be driven. The results are below:

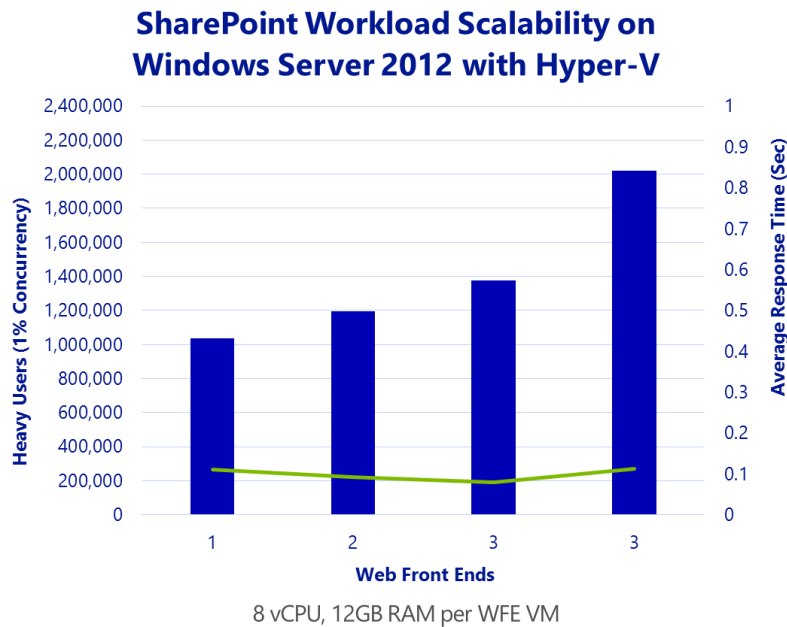


Figure 25 – Graph of virtualized SharePoint 2013 scalability with increased demand

Testing found that 3 WFEs could support just over 2 million heavy users at 1% concurrency, with an average CPU utilization of 84% across WFEs, with measurably low response times.

#### Why This Matters

As customers move past virtualizing their less critical systems, and move onto the more powerful, more demanding, more mission-critical systems, their confidence that virtualization can provide what they had previously in the physical world is of utmost importance. Performance is one of the key considerations to virtualizing these kind of workloads, and with the performance testing by ESG, along with the key scale and performance capabilities in Windows Server 2012 R2 Hyper-V, it's clear that customers can successfully virtualize those workloads, without sacrifices, on a cost-effective virtualization platform.

# Security & Multitenancy

Virtualized data centers are becoming more popular and practical every day. IT organizations and hosting providers have begun offering infrastructure as a service (IaaS), which provides more flexible, virtualized infrastructures to customers—“server instances on-demand.” Because of this trend, IT organizations and hosting providers must offer customers enhanced security and isolation from one another.

If a service provider’s infrastructure is hosting two companies, the IT Admin must help ensure that each company is provided its own privacy and security. Before Windows Server 2012 and subsequently, Windows Server 2012 R2, server virtualization provided isolation between virtual machines, but the network layer of the data center was still not fully isolated and implied layer-2 connectivity between different workloads that run over the same infrastructure.

For the hosting provider, isolation in the virtualized environment must be equal to isolation in the physical data center, to meet customer expectations and not be a barrier to cloud adoption.

Isolation is almost as important in an enterprise environment. Although all internal departments belong to the same organization, certain workloads and environments (such as finance and human resource systems) must still be isolated from each other. IT departments that offer private clouds and move to an IaaS operational mode must consider this requirement and provide a way to isolate such highly sensitive workloads.

Windows Server 2012 R2 contains powerful and comprehensive security and isolation capabilities that are provided as part of the Hyper-V Extensible Switch.

## The Hyper-V Extensible Switch

The Hyper-V Extensible Switch is a layer-2 virtual network switch that provides programmatically managed and extensible capabilities to connect virtual machines to the physical network with policy enforcement for security and isolation. The figure below shows a network using the Hyper-V Extensible Switch.

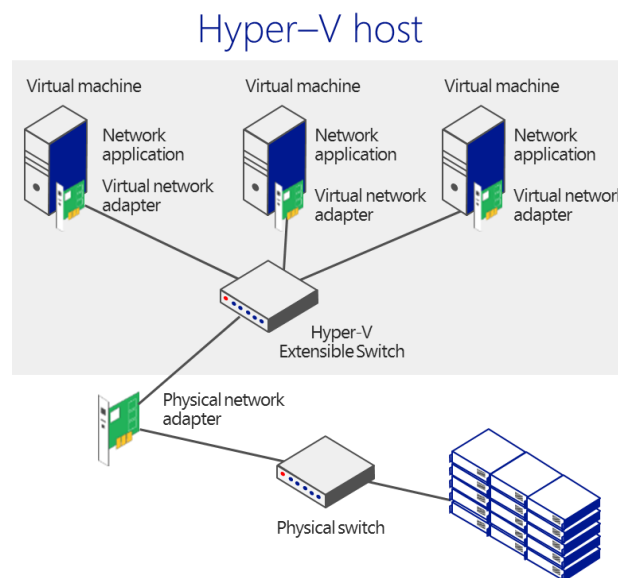


Figure 26 – Network with Hyper-V Extensible Switch

With Windows Server 2012 R2, the IT Admin can configure Hyper-V servers to enforce network isolation among any set of arbitrary isolation groups, which are typically defined for individual customers or sets of workloads. Windows Server 2012 R2 provides the isolation and security capabilities for multitenancy by offering the following key features:

- **Private VLANs (PVLANS)** - Provide isolation between two virtual machines on the same VLAN
- **ARP/ND Poisoning/Spoofing** - Protection against a malicious virtual machine stealing IP addresses from other virtual machines
- **DHCP Snooping/DHCP Guard** - Protects against rogue DHCP servers attempting to provide IP addresses that would cause traffic to be rerouted
- **Virtual Port ACLs** - Isolate networks and metering network traffic for a virtual port
- **Trunk Mode to Virtual Machines** - Traffic from multiple VLANs can now be directed to a single network adapter in a virtual machine
- **Monitoring & Port Mirroring** - Monitor the traffic from specific ports flowing through specific virtual machines on the switch and mirror traffic which can then be delivered to another virtual port for further processing
- **Windows PowerShell/Windows Management Instrumentation (WMI)** - Provides Windows PowerShell cmdlets for the Hyper-V Extensible Switch that lets customers and partners build command-line tools or automated scripts for setup, configuration, monitoring, and troubleshooting

## PVLANS

VLAN technology is traditionally used to subdivide a network and provide isolation for individual groups that share a common physical infrastructure. Windows Server 2012 R2 supports PVLANS, a capability used with VLANs that can be used to provide isolation between two virtual machines on the same VLAN.

When a virtual machine doesn't need to communicate with other virtual machines, you can use PVLANS to isolate it from other virtual machines in your datacenter. By assigning each virtual machine in a PVLAN, one primary VLAN ID and one or more secondary VLAN IDs, you can put the secondary PVLANS into one of three modes (as shown in the following table). These PVLAN modes determine which other virtual machines on the PVLAN a virtual machine can talk to. To isolate a virtual machine, put it in isolated mode.

PVLAN Mode	Description
Isolated	Isolated ports cannot exchange packets with each other at layer 2.
Promiscuous	Promiscuous ports can exchange packets with any other port on the same primary VLAN ID.
Community	Community ports on the same VLAN ID can exchange packets with each other at layer 2.

Table 3 – PVLAN modes for virtual machine isolation

The following figure shows how the three PVLAN modes can be used to isolate virtual machines that share a primary VLAN ID.

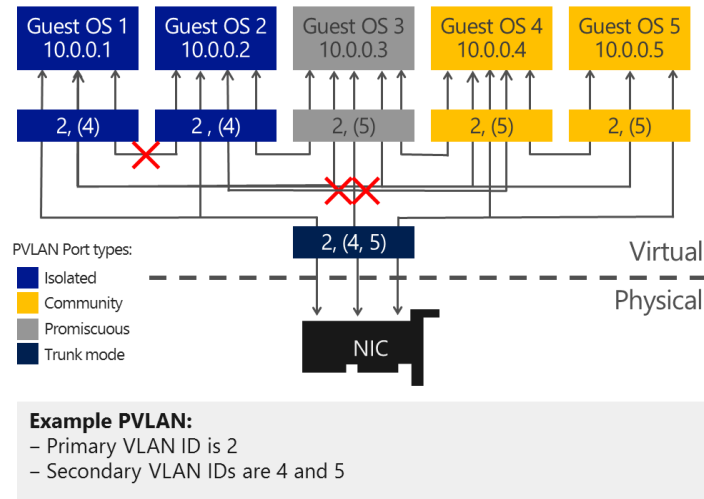


Figure 27 – Example PVLAN with primary VLAN ID 2

## ARP/ND Poisoning and Spoofing Protection

The Hyper-V Extensible Switch provides protection against a malicious virtual machine stealing IP addresses from other virtual machines through ARP spoofing (also known as ARP poisoning in IPv4). With this type of man-in-the-middle attack, a malicious virtual machine sends a fake ARP message, which associates its own MAC address to an IP address that it doesn't own. Unsuspecting virtual machines send network traffic targeted to that IP address to the MAC address of the malicious virtual machine instead of the intended destination. For IPv6, Windows Server 2012 provides equivalent protection for ND spoofing.

## DHCP Guard Protection

In a DHCP environment, a rogue DHCP server could intercept client DHCP requests and provide incorrect address information. The rogue DHCP server could cause traffic to be routed to a malicious intermediary that sniffs all traffic before forwarding it to the legitimate destination. To protect against this particular man-in-the-middle attack, the Hyper-V administrator can designate which Hyper-V Extensible Switch ports can have DHCP servers connected to them. DHCP server traffic from other Hyper-V Extensible Switch ports is automatically dropped. The Hyper-V Extensible Switch now protects against a rogue DHCP server attempting to provide IP addresses that would cause traffic to be rerouted.

## Hyper-V Virtual Switch Extended Port ACLs

Enterprises and Cloud Service Providers (CSPs) can configure the Hyper-V Virtual Switch Extended Port Access Control Lists (ACLs) to provide firewall protection and enforce security policies for the tenant VMs in their datacenters. Because the port ACLs are configured on the Hyper-V Virtual Switch rather than within the VMs, you can manage security policies for all tenants in a multitenant environment

By using port ACLs, you can meter the IP addresses or MAC addresses that can (or cannot) communicate with a virtual machine. For example, you can use port ACLs to enforce isolation of a virtual machine by letting it talk only to the Internet, or communicate only with a predefined set of addresses. By using the metering capability, you can measure network traffic going to or from a specific IP address or MAC address, which lets you report on traffic, sent or received from the Internet or from network storage arrays.

You also can configure multiple port ACLs for a virtual port. Each port ACL consists of a source or destination network address, and a permit to deny or meter action. The metering capability also supplies information about the number of instances where traffic was attempted to or from a virtual machine from a restricted ("deny") address.

Port ACLs were first introduced in Windows Server 2012 Hyper-V, however there have been a number of key improvements introduced as part of Windows Server 2012 R2 Hyper-V:

- ACLs now include the socket port number. In Windows Server 2012, you were able to specify both source and destination MAC and IP addresses for IPv4 and IPv6. For Windows Server 2012 R2 you can also specify the port number when you create rules
- You can now configure stateful rules that are unidirectional and provide a timeout parameter. With a stateful firewall rule, traffic is allowed, and two traffic flows are created dynamically. The two traffic flows are one outbound rule that matches five attributes in outbound packets, and one inbound rule that also match the same five attributes. After a stateful rule is utilized successfully one time, the two traffic flows are allowed without having to be looked up against the rule again for a period of time that you designate using the timeout attribute. When the firewall rule exceeds the timeout attribute, traffic flows are inspected against rules again.

In addition, extended port ACLs provide the following benefits:

- In multitenant environments, you can protect datacenter resources and provide security policy enforcement for your tenants.
- Compatibility with Hyper-V Network Virtualization.
- A management interface that allows you to easily configure firewall rules by using Windows PowerShell.
- Logging and diagnostics capabilities so that you can confirm firewall operation and detect any possible misconfiguration of the port ACLs.
- Configurable as a stateless firewall by filtering packets based on five attributes in the packet; with a stateless firewall configuration you can apply any firewall rule to either inbound or outbound network traffic, and the rule can either allow or deny traffic

## Trunk Mode to Virtual Machines

A VLAN makes a set of host machines or virtual machines appear to be on the same local LAN, independent of their actual physical locations. With the Hyper-V Extensible Switch trunk mode, traffic from multiple VLANs can now be directed to a single network adapter in a virtual machine that could previously receive traffic from only one VLAN. As a result, traffic from different VLANs is consolidated, and a virtual machine can listen in on multiple VLANs. This feature can help you shape network traffic and enforce multitenant security in your datacenter.

## Monitoring

Many physical switches can monitor the traffic from specific ports flowing through specific virtual machines on the switch. The Hyper-V Extensible Switch also provides this port mirroring, enabling you to designate which virtual ports should be monitored and to which virtual port the monitored traffic should be delivered for further processing. For example, a security-monitoring virtual machine can look for anomalous patterns in the traffic that flows through other specific virtual machines on the switch. In addition, you can diagnose network connectivity issues by monitoring traffic bound for a particular virtual switch port.

## Windows PowerShell and WMI

Windows Server 2012 R2 provides Windows PowerShell cmdlets for the Hyper-V Extensible Switch that lets you build command-line tools or automated scripts for setup, configuration, monitoring, and troubleshooting. These cmdlets can be run remotely. Windows PowerShell also enables third parties to build their own tools to manage the Hyper-V Extensible Switch.

## Requirements

To take advantage of the Hyper-V Extensible Switch, and all of the key capabilities discussed above, you will require the following:

- Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2

### Why This Matters

**Windows Server 2012 R2 multitenant isolation keeps customer virtual machines isolated, even when they are stored on the same physical server. Windows Server 2012 R2 provides better multitenant security for customers on a shared IaaS cloud through the new Hyper-V Extensible Switch, which provides:**

- **Security and isolation.** The Hyper-V Extensible Switch provides better security and isolation for IaaS multitenancy with PVLAN support, protection against ARP poisoning and spoofing, protection against DHCP snooping, virtual port ACLs, and VLAN trunk mode support.
- **Monitoring.** With port mirroring, you can run security and diagnostics applications in virtual machines that can monitor virtual machine network traffic. Port mirroring also supports live migration of extension configurations.
- **Manageability.** You can now use Windows PowerShell and WMI support for command-line and automated scripting support, as well as full event logging.

**Multitenant isolation in Windows Server 2012 R2 addresses concerns that may have previously prevented organizations from deploying Hyper-V within the datacenters. Two such concerns are:**

- **Additional management overhead of implementing VLANs on the Ethernet switching infrastructure to ensure isolation between their customers' virtual infrastructures.**
- **Security risk of a multitenant virtualized environment.**

**With Hyper-V in Windows Server 2012 R2, you can use port ACLs to isolate customers' networks from one another and not be required to set up and maintain VLANs. Also, your security needs are met by protection against ARP spoofing and DHCP snooping.**

## Extending the Extensible Switch

Many enterprises need the ability to extend virtual switch features with their own plug-ins to suit their virtual environment. When IT professionals install virtual switches, they naturally look for the same kind of functionality that they can achieve on physical networks, such as adding firewalls, intrusion detection systems, and network traffic monitoring tools. However, the challenge has been finding easy ways to add virtualized appliances, extensions, and other features and functions to virtual switches. Most virtual switch technology offerings are built around closed systems that make it difficult for enterprise developers and third-party vendors to build solutions and to quickly and easily install new functionality into their virtual switches.

The Hyper-V Extensible Switch changes all that. With the Hyper-V Extensible Switch, IT professionals can easily add more functionality to their virtual machines and networks. At the same time, it gives internal enterprise developers and third-party providers an open platform for creating solutions that extend the basic functionality of the switch. If you're in charge of making IT purchasing decisions at your company, you want to know that the virtualization platform you choose won't lock you in to a small set of compatible features, devices, or technologies. In Windows Server 2012 R2, the Hyper-V Extensible Switch provides key extensibility features.

The Hyper-V Extensible Switch is an open platform that lets multiple vendors provide extensions that are written to standard Windows API frameworks. The reliability of extensions is strengthened through the Windows standard framework and reduction of required third-party code for functions and is backed by the Windows Hardware Quality Labs (WHQL) certification program. The IT Admin can manage the Hyper-V Extensible Switch and its extensions by using Windows PowerShell, programmatically with WMI, through the Hyper-V Manager user interface, or through System Center Virtual Machine Manager 2012 R2.

## Extensibility

The Hyper-V Extensible Switch architecture in Windows Server 2012 R2 is an open framework that lets third parties add new functionality such as monitoring, forwarding, and filtering into the virtual switch. Extensions are implemented by using Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers. These two public Windows platforms for extending Windows networking functionality are used as follows:

- **NDIS filter drivers** are used to monitor or modify network packets in Windows. NDIS filters were introduced with the [NDIS 6.0 specification](#).
- **WFP callout drivers**, introduced in Windows Vista and Windows Server 2008, let independent software vendors (ISVs) create drivers to filter and modify TCP/IP packets, monitor or authorize connections, filter IP security (IPsec)-protected traffic, and filter remote procedure calls (RPCs). Filtering and modifying TCP/IP packets provides unprecedented access to the TCP/IP packet processing path. In this path, you can examine or modify outgoing and incoming packets before additional processing occurs. By accessing the TCP/IP processing path at different layers, you can more easily create firewalls, antivirus software, diagnostic software, and other types of applications and services. For more information, see the [Windows Filtering Platform](#).

Extensions may extend or replace three aspects of the switching process:

- Ingress filtering.
- Destination lookup and forwarding.
- Egress filtering.

In addition, by monitoring extensions you can gather statistical data by monitoring traffic at different layers of the switch. Multiple monitoring and filtering extensions can be supported at the ingress and egress portions of the Hyper-V Extensible Switch. Only one instance of the forwarding extension may be used per switch instance, and it overrides the default switching of the Hyper-V Extensible Switch.

The table below, lists the various types of Hyper-V Extensible Switch extensions.

Extension	Purpose	Examples	Extensibility Component
Network Packet Inspection	Inspecting network packets, but not altering them	sFlow and network monitoring	NDIS filter driver

Extension	Purpose	Examples	Extensibility Component
Network Packet Filtering	Injecting, modifying, and dropping network packets.	Security	NDIS filter driver
Network Forwarding	Third-party forwarding that bypasses default forwarding	OpenFlow, Virtual Ethernet Port Aggregator (VEPA), and proprietary network fabrics	NDIS filter driver
Firewall/Intrusion Detection	Filtering and modifying TCP/IP packets, monitoring or authorizing connections, filtering IPsec-protected traffic, and filtering RPCs.	Virtual firewall and connection monitoring	WFP callout driver

Table 4 – Types of Hyper-V Extensible Switch extensions

The Hyper-V Extensible Switch provides an open-switch API that lets enhanced switch and management products work with Hyper-V.

The Hyper-V Extensible Switch architecture in Windows Server 2012 R2 is an open framework that lets third parties add new functionality into the virtual switch. The following figure shows the architecture of the Hyper-V Extensible Switch and the extensibility model.

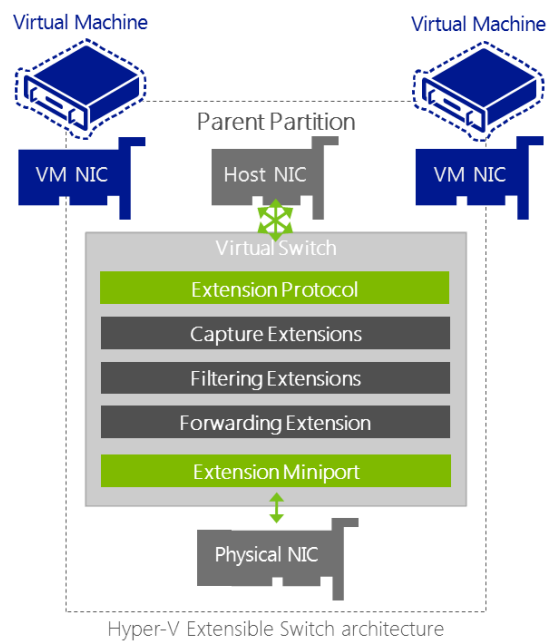


Figure 28 – Architecture of the Hyper-V Extensible Switch

Some other features of Hyper-V Extensible Switch extensibility are:

- **Extension monitoring.** Monitoring extensions lets you gather statistical data by monitoring traffic at different layers of the Hyper-V Extensible Switch. Multiple monitoring and filtering extensions can be supported at the ingress and egress portions of the Hyper-V Extensible Switch.
- **Extension uniqueness.** Extension state/configuration is unique to each instance of an Extensible Switch on a machine.
- **Extensions that learn from virtual machine life cycle.** A virtual machine's activity cycle is similar to that of physical servers, having peak times during various times of the day or night based on its core



workloads. Extensions can learn the flow of network traffic based on the workload cycles of your virtual machines, and optimize your virtual network for greater performance.

- **Extensions that can veto state changes.** Extensions can implement monitoring, security, and other features to further improve the performance, management, and diagnostic enhancements of the Hyper-V Extensible Switch. Extensions can help ensure the security and reliability of your system by identifying and blocking implementation of harmful state changes.
- **Multiple extensions on same switch.** Multiple extensions can coexist on the same Hyper-V Extensible Switch.

## Manageability

By using the following management features built into the Hyper-V Extensible Switch, you can troubleshoot and resolve problems on Hyper-V Extensible Switch networks:

- **Windows PowerShell and scripting support.** Windows Server 2012 R2 provides Windows PowerShell cmdlets for the Hyper-V Extensible Switch that let you build command-line tools or automated scripts for setup, configuration, monitoring, and troubleshooting. Windows PowerShell also enables third parties to build their own Windows PowerShell-based tools to manage the Hyper-V Extensible Switch.
- **Unified tracing and enhanced diagnostics.** The Hyper-V Extensible Switch includes unified tracing to provide two levels of troubleshooting. At the first level, the Event Tracing for Windows (ETW) provider for the Hyper-V Extensible Switch permits tracing packet events through the Hyper-V Extensible Switch and extensions, making it easier to pinpoint where an issue occurred. The second level permits capturing packets for a full trace of events and traffic packets.

## Partner Extensions

As mentioned earlier, many enterprises need the ability to extend virtual switch features with their own plug-ins to suit their virtual environment, often to mimic or replicate what they had within the physical environment. With the Hyper-V Extensible Switch, and a rapidly growing Partner ecosystem, customers can integrate, or even build specific functionality on top of the core vSwitch to enable new scenarios specific to their needs.

Several Partners have already announced, and have released extensions for the Hyper-V Extensible Switch, including:

- **Cisco - Nexus 1000V Series Switch & UCS Virtual Machine Fabric Extender (VM-FEX).** The **Cisco Nexus 1000V Switch** offers a consistent operational model across physical and virtual environments. This distributed virtual switching platform provides advanced features and is tightly integrated with the Hyper-V ecosystem. **Cisco Virtual Machine Fabric Extender (VM-FEX)** collapses virtual and physical networking into a single infrastructure. Data center administrators can now provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare metal network traffic within a unified infrastructure.
- **NEC – PF1000.** The **ProgrammableFlow PF1000 virtual switch** simplifies complex networks and integrates server and network virtualization within a single control pane, and brings supports for OpenFlow to Hyper-V.
- **Snine – Security Manager** provides an agentless anti-virus and anti-malware protection, along with a powerful Kernel Mode virtual firewall that delivers comprehensive real-time traffic filtering for the virtualized environment
- **InMon – sFlow.** The **sFlow** standard provides an integrated end-to-end view of performance. It defines a coherent framework of metrics for integrated network, server and application performance monitoring.

## Requirements

To take advantage of the Hyper-V Extensible Switch, and all of the key capabilities discussed earlier, along with the ability to plug in Partner extensions, you will require the following:

- Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2

### Why This Matters

**The Hyper-V Extensible Switch is an open platform, so third-party vendors can provide plug-ins that supply additional functionality such as traffic monitoring, firewall filters, and switch forwarding. Plug-in management is unified through Windows PowerShell cmdlets and WMI scripting.**

**The Hyper-V Extensible Switch permits easier implementation and management of virtualized datacenters by providing the following:**

- **Open platform to fuel plug-ins.** The Hyper-V Extensible Switch is an open platform that lets plug-ins sit in the virtual switch between all traffic, including virtual machine-to-virtual machine traffic. Extensions can provide traffic monitoring, firewall filters, and switch forwarding. To jump-start the ecosystem, several partners will announce extensions when the Hyper-V Extensible Switch is released. No “one-switch-only” solution for Hyper-V will occur.
- **Core services provided at no cost.** Core services are provided for extensions. For example, all extensions have live migration support by default, and no special coding for services is required.
- **Windows reliability and quality.** Extensions provide a high level of reliability and quality from the strength of the Windows platform and the Windows logo certification program, both of which set a high bar for extension quality.
- **Unified management.** Managing extensions is integrated into Windows management through Windows PowerShell cmdlets and WMI scripting.
- **Easier support.** Unified tracing makes it quicker and easier to diagnose any issues that arise. This means less downtime and increased availability of services.
- **Live migration support.** The Hyper-V Extensible Switch provides capabilities enabling extensions to participate in Hyper-V live migration.

**The Hyper-V Extensible Switch gives third-party vendors the freedom to develop custom solutions for handling network traffic in a Windows Server 2012 R2 virtual network. For example, these solutions can be used to emulate a vendor’s physical switch and its policies, or to monitor and analyze traffic.**

## Physical Security

When it comes to deployment of virtualization technologies, many are within secure datacenter environments, but what about those that aren’t? Satellite offices, remote sites, home offices and retail stores are all examples of environments that may not have them same levels of physical security as the enterprise datacenter, yet may still have physical servers, with virtualization technologies present. If the

physical hosts were compromised, there could be very serious repercussions for the business. What if there are compliancy requirements, to have an encrypted environment?

## BitLocker

With **Windows Server 2012 R2 Hyper-V**, BitLocker Drive Encryption is included to solve those very problems, by allowing customers to encrypt all data stored on the Windows Server 2012 R2 operating system volume and configured data volumes, along with any Failover Cluster disks, including Cluster Shared Volumes, ensuring that environments, large and small, that are implemented in less physically secure locations, can have the highest levels of data protection for their key workloads, at no additional cost.




Figure 29 – Supported locations with BitLocker Drive Encryption

As you can see from the figure above, in Windows Server 2012 R2 Hyper-V, Microsoft provides support for encrypting the following repositories that are used to store virtual machine disks:

- Local Disk i.e. DAS
- A File Server Share (BitLocker would be enabled on the File Server)
- A traditional LUN which has been made available to the Hyper-V host, and formatted appropriately.
- A Cluster Shared Volume

BitLocker in Windows Server 2012 R2 has a number of useful capabilities specific to the server environment that IT administrators can benefit from:

- **Used Disk Space Only encryption** - BitLocker now offers two encryption methods, Used Disk Space Only and Full volume encryption. Used Disk Space Only allows for a much quicker encryption experience by only encrypting used blocks on the targeted volume.
- **Network Unlock** - Enables a BitLocker system on a wired network to automatically unlock the system volume during boot/reboot (on capable Windows Server 2012 networks), reducing internal help desk call volumes for lost PINs. This is incredibly important in a lights-out datacenter.

## Requirements

To take advantage of BitLocker and Hyper-V Virtual Machines, you will require the following:

- Windows Server 2012 with Hyper-V, Windows Server 2012 R2 with Hyper-V, Hyper-V Server 2012, or **Hyper-V Server 2012 R2**

### Why This Matters

Customers with compliancy needs that require all data to be encrypted, or have physical servers located in less-physically secure locations, can benefit considerably from the built-in capabilities that BitLocker provides. Simple to enable, and easy to manage, BitLocker is a powerful encryption technology that helps to keep your data safe, and secure, across a multitude of different storage options. Once enabled, BitLocker is transparent to ongoing operation, VM management and deployment, integrates with AD for management, and ensures that in the event of a physical compromise, such as a stolen server, that data residing on the host, and inside the VMs, is secure.

## Flexible Infrastructure

We've spent a considerable amount of time discussing some of the key capabilities within Windows Server 2012 R2, and how they combine to provide huge scalability, performance and density, to run the most mission critical of workloads, and in addition, how the security capabilities provided by features such as BitLocker, and those ingrained within the Hyper-V Extensible Switch, ensure that those workloads are secure, and isolated against malicious attacks. These however, aren't the only important considerations customers must make when evaluating a virtualization platform. Customers decide to virtualize their workloads not just for consolidation's sake, but to gain new levels of flexibility and agility. This could be in the form of having greater flexibility to move workloads around the infrastructure to best utilize existing hardware capacity, or to mitigate hardware maintenance. Alternatively, it could be flexibility to deploy new workloads, and manage their placement within isolated virtualized networks. It could also be the flexibility to run workloads that aren't Windows-based, such as those running on Linux. Fortunately, Hyper-V in Windows Server 2012 R2 provides solutions for each of these areas.

### Linux Support on Hyper-V

The ability to provision Linux on Hyper-V and Windows Azure is one of Microsoft's core efforts towards enabling great Open Source Software support. As part of this initiative, the **Microsoft Linux Integration Services (LIS)** team pursues ongoing development of enlightened Linux drivers that are directly checked in to the Linux upstream kernel thereby allowing direct integration into upcoming releases of major distributions such as CentOS, Debian, Red Hat, SUSE and Ubuntu.

The Integration Services were originally shipped as a download from Microsoft's sites. Linux users could download and install these drivers and contact Microsoft for any requisite support. As the drivers have matured, they are now delivered directly through the Linux distributions. Not only does this approach avoid the extra step of downloading drivers from Microsoft's site but it also allows users to leverage their existing support contracts with Linux vendors.

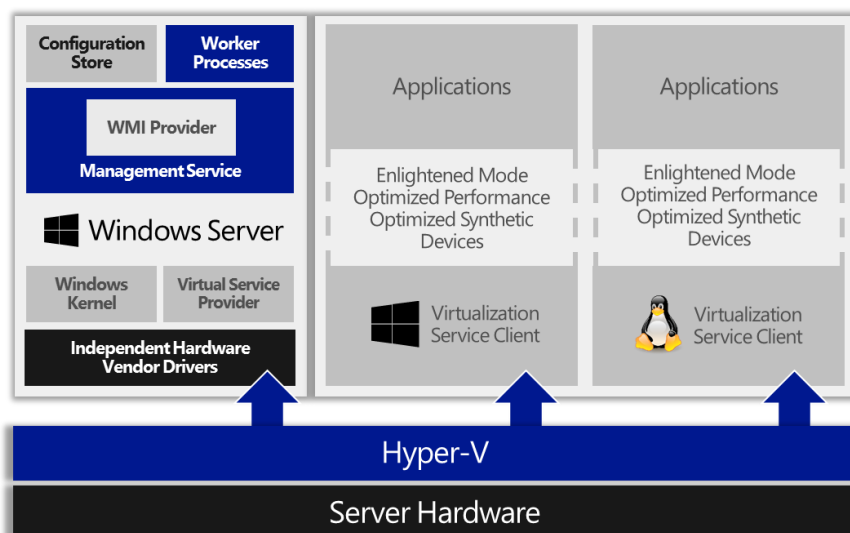


Figure 30 – Architectural view of Linux running enlightened on Hyper-V

For example Red Hat has certified enlightened drivers for Hyper-V on Red Hat Enterprise Linux (RHEL) 5.9 and certification of RHEL 6.4 should be complete by summer 2013. This will allow customers to directly obtain Red Hat support for any issues encountered while running RHEL 5.9/6.4 on Hyper-V.

To further the goal of providing great functionality and performance for Linux running on Microsoft infrastructure, there are a number of new features in Windows Server 2012 R2 that are now available for Linux guest operating systems.

## 64 vCPU Support & Deeper Integration Services Support

With the launch of Windows Server 2012, and subsequently, Windows Server 2012 R2, Hyper-V provides support for virtual machines running Linux, with up to 64 virtual processors, driving incredible scalability and performance for CPU-intensive workloads. In addition, there have been significant improvements to the Linux Integration Services to drive improved performance for Network and Storage controllers, Fastpath Boot Support, Time Keeping, Integrated shutdown and more.

## Synthetic Frame Buffer Driver

The new synthetic 2D frame buffer driver provides solid improvements in graphics performance for Linux virtual machines running on Hyper-V. Furthermore, the driver provides full HD mode resolution (1920x1080) capabilities for Linux guests hosted in desktop mode on Hyper-V.

One other noticeable impact of the Synthetic Frame Buffer Driver is elimination of the double cursor problem. While using desktop mode on older Linux distributions several customers reported two visible mouse pointers that appeared to chase each other on screen. This distracting issue is now resolved through the synthetic 2D frame buffer driver thereby improving visual experience on Linux desktop users.

## Dynamic Memory Support

The availability of dynamic memory for Linux guests provides higher virtual machine density per host. This will bring huge value to Linux administrators looking to consolidate their server workloads using Hyper-V. In house test results indicate a 30-40% increase in server capacity when running Linux machines configured with dynamic memory.

The Linux dynamic memory driver monitors the memory usage within a Linux virtual machine and reports it back to Hyper-V on a periodic basis. Based on the usage reports Hyper-V dynamically orchestrates memory allocation and deallocation across various virtual machines being hosted. Note that the user interface for configuring dynamic memory is the same for both Linux and Windows virtual machines.

The dynamic Memory driver for Linux virtual machines provides both Hot-Add and Ballooning support and can be configured using the Start, Minimum RAM and Maximum RAM parameters in Hyper-V Manager.

Upon system start the Linux virtual machine is booted up with the amount of memory specified in the Start parameter. If the virtual machine requires more memory then Hyper-V uses the Hot-Add mechanism to dynamically increase the amount of memory available to the virtual machine.

On the other hand, if the virtual machine requires less memory than allocated then Hyper-V uses the ballooning mechanism to reduce the memory available to the virtual machine to a more appropriate amount.

## Live Virtual Machine Backup Support

A much requested feature from customers running Linux on Hyper-V is the ability to create seamless backups of live Linux virtual machines. In the past customers had to either suspend or shutdown the Linux virtual machine for creating backups. Not only is this process hard to automate but it also leads to an increase in down time for critical workloads.

To address this shortcoming, a file-system snapshot driver is now available for Linux guests running on Hyper-V. Standard backup APIs available on Hyper-V can be used to trigger the driver to create file-system

consistent snapshots of VHDs attached to a Linux virtual machine without disrupting any operations in execution within the virtual machine.

One important difference between the backups of Linux virtual machines and Windows virtual machines is that Linux backups are file-system consistent only whereas Windows backups are file-system and application consistent. This difference is due to lack of standardized Volume Shadow Copy Service (VSS) infrastructure in Linux.

### Dynamic Expansion of Live Fixed Sized VHDXs

The ability to dynamically resize a VHDX allows administrators to allocate more storage to the VHD while keeping the performance benefits of the format. The feature is now available for Linux virtual machines running on Hyper-V. It is worth noting that Linux file-systems are quite adaptable to dynamic changes in size of the underlying disk drive.

### Linux kdump/kexec support

One particular pain point for hosters running Linux on Windows Server 2012 and Windows Server 2008 R2 environments is that legacy drivers (as mentioned in [KB 2858695](#)) must be used to create kernel dumps for Linux virtual machines.

In Windows Server 2012 R2, the Hyper-V infrastructure has been changed to allow seamless creation of crash dumps using enlightened storage and network drivers and therefore no special configurations are required anymore. Linux users are free to dump core over the network or the attached storage devices.

### NMI Support

If a Linux system becomes completely unresponsive while running on Hyper-V, users now have the option to panic the system by using Non-Maskable Interrupts (NMI). This is particularly useful for diagnosing systems that have deadlocked due to kernel or user mode components.

### Specification of Memory Mapped I/O (MMIO) Gap

Linux based appliance manufacturers use the MMIO gap (also known as PCI hole) to divide the available physical memory between the Just Enough Operating System (JeOS) that boots up the appliance and the actual software infrastructure that powers the appliance. Inability to configure the MMIO gap causes the JeOS to consume all of the available memory leaving nothing for the appliance's custom software infrastructure. This shortcoming inhibits the development of Hyper-V based virtual appliances.

The Windows Server 2012 R2 Hyper-V infrastructure allows appliance manufacturers to configure the location of the MMIO gap. Availability of this feature facilitates the provisioning of Hyper-V powered virtual appliances in hosted environments.

### Requirements

To take advantage of the features listed above, with your Linux distributions, you will require the following:

- Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2
- **For Linux distributions without the LIS drivers built in, you will need to download the current LIS drivers from the [Microsoft.com download site](#).**

## Why This Matters

Over the past year, the LIS team added a slew of features to enable great support for Linux virtual machines running on Hyper-V. These features will not only simplify the process of hosting Linux on Hyper-V but will also provide superior consolidation and improved performance for Linux workloads. This provides customers, and partners, with the confidence to deploy not just Microsoft workloads, but also workloads based on Linux, and in addition, with our relationships with both the Linux community, and Linux commercial vendors, the platform and integration is growing, release on release, so customers can expect near-parity with virtualized Windows workloads in the future.

## Virtual Machine Live Cloning

When customers have virtual machines with applications running inside, and those application experiences issues, customers have a difficult choice to make. Do they leave the application up and running, and ask IT to create and deploy a workload, from scratch, to match this workload, in order to try to replicate the issue exactly? Do they shut the workload down to reduce the risk of further issues, so that IT can fix the application without the applications being used by end users? In Windows Server 2012 Hyper-V, customers would need to shut down the virtual machine, and then clone, meaning there would be downtime for the application, however with Windows Server 2012 R2, things have changed considerably.

Windows Server 2012 R2 now supports a feature known as Live Cloning.




Figure 31 – Virtual Machine Live Cloning initialized and cloning the key files

The IT administrator initializes the Live Clone, or technically, a Live Export, whilst the VM is running, and Hyper-V will export the appropriate files to the new target location, all whilst the source VM continues to run.





Figure 32 – New VM is imported into Hyper-V Manager

In addition to being able to clone a running VM, any VMs that currently have checkpoints (snapshots) created, can also be exported and cloned with no downtime. The IT Administrator can even create a complete new VM from a specific checkpoint. This could be very useful for troubleshooting an issue that occurred as a result of a software configuration change or untested patch for example.

### Requirements

To take advantage of Live Cloning capability, you will require the following:

- Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2

### Why This Matters

**Problems happen. Whether down to misconfiguration within an application, an untested patch being deployed, or simply running a previously untested scenario, when the problem occurs, in order to test and replicate as quickly as possible, IT administrators can use the new Live Cloning capability within Hyper-V. This feature enables the IT administrator to clone the VM, without taking the existing VM down, and provide that VM to the development or application teams in order to troubleshoot the issue on a system which matches the original identically. This will enable a faster time to resolution, and simplified maintenance.**

## Virtual Machine Mobility

To maintain optimal use of physical resources and to be able to easily add new virtual machines, IT must be able to move virtual machines whenever necessary without disrupting the business. The ability to move virtual machines across Hyper-V hosts was introduced in Windows Server 2008 R2, as a feature known as Live Migration, however this was limited to the VM residing on shared storage, being part of a cluster, and, multiple VMs couldn't be moved simultaneously. With Windows Server 2012, and subsequently, Windows Server 2012 R2, Hyper-V has a number of new and powerful migration capabilities that provide complete flexibility for customers wishing to migrate workloads around their datacenter.

### Live Migration

Hyper-V live migration moves running virtual machines from one physical server to another with no impact on virtual machine availability to users. By pre-copying the memory of the migrating virtual machine to the

destination server, live migration minimizes the transfer time of the virtual machine. A live migration is deterministic, which means that the administrator, or script, that initiates the live migration determines which computer is used as the destination for the live migration. The guest operating system of the migrating virtual machine is not aware that the migration is happening, so no special configuration for the guest operating system is needed.

After initiating a live migration, the following process occurs:

1. **Live migration setup occurs** - During the live migration setup stage, the source server creates a connection with the destination server. This connection transfers the virtual machine configuration data to the destination server. A skeleton virtual machine is set up on the destination server and memory is allocated to the destination virtual machine.
2. **Memory pages are transferred from the source node to the destination node** - In the second stage of a live migration, the memory assigned to the migrating virtual machine is copied over the network to the destination server. This memory is referred to as the "working set" of the migrating virtual machine. A page of memory is 4 KB.

For example, suppose that a virtual machine named "test virtual machine" configured with 1024 MB of RAM is migrating to another server running Hyper-V. The entire 1024 MB of RAM assigned to this virtual machine is the working set of "test virtual machine." The utilized pages within the "test virtual machine" working set are copied to the destination server.

In addition to copying the working set of "test virtual machine" to the destination server, Hyper-V monitors the pages in the working set for "test virtual machine" on the source server. As memory pages are modified by "test virtual machine," they are tracked and marked as being modified. The list of modified pages is simply the list of memory pages "test virtual machine" has modified after the copy of its working set has begun.

During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, with each iteration requiring a smaller number of modified pages to be copied. After the working set is copied to the destination server, the next stage of the live migration begins.

3. **Modified pages are transferred** - The third stage of a live migration is a memory copy process that duplicates the remaining modified memory pages for "test virtual machine" to the destination server. The source server transfers the CPU and device state of the virtual machine to the destination server.

During this stage, the network bandwidth available between the source and destination servers is critical to the speed of the live migration. Using a 1 Gigabit Ethernet or faster is important. The faster the source server transfers the modified pages from the migrating virtual machines working set, the more quickly the live migration is completed.

The number of pages transferred in this stage is determined by how actively the virtual machine accesses and modifies the memory pages. The more modified pages there are, the longer it takes to transfer all pages to the destination server.

After the modified memory pages are copied completely to the destination server, the destination server has an up-to-date working set for "test virtual machine." The working set for "test virtual machine" is present on the destination server in the exact state it was in when "test virtual machine" began the migration process.

4. **The storage handle is moved from the source server to the destination server** - During the fourth stage of a live migration, control of the storage associated with "test virtual machine," such as any virtual hard disk files or physical storage attached through a virtual Fibre Channel adapter, is transferred to the destination server.

5. **The virtual machine is brought online on the destination server** - In the fifth stage of a live migration, the destination server now has the up-to-date working set for “test virtual machine,” as well as access to any storage used by “test virtual machine.” At this point “test virtual machine” is resumed.
6. **Network cleanup occurs** - In the final stage of a live migration, the migrated virtual machine is running on the destination server. At this point, a message is sent to the network switch. This message causes the network switch to obtain the new the MAC addresses of the migrated virtual machine so that network traffic to and from “test virtual machine” can use the correct switch port.

The live migration process completes in less time than the TCP time-out interval for the virtual machine being migrated. TCP time-out intervals vary based on network topology and other factors. The following variables may affect live migration speed:

- The number of modified pages on the virtual machine to be migrated—the larger the number of modified pages, the longer the virtual machine will remain in a migrating state.
- Available network bandwidth between source and destination servers.
- Hardware configuration of source and destination servers.
- Load on source and destination servers.
- Available bandwidth (network or Fibre Channel) between servers running Hyper-V and shared storage.

## SMB-Based Live Migration

In Windows Server 2012 R2 Hyper-V, you can configure a virtual machine so that it is stored on an SMB file share. You can then perform a live migration on this running virtual machine between non-clustered servers running Hyper-V, while the virtual machine’s storage remains on the central SMB share. This allows users to gain the benefits of virtual machine mobility without having to invest in the clustering infrastructure if they do not need guarantees of availability in their environment. (Hyper-V with SMB storage can also be configured with Failover Clustering if you do require high availability).

## Faster and Simultaneous Migrations

If you use live migration in a clustered environment today, you will see that live migrations can now use higher network bandwidths (up to 10 GbE) to complete migrations faster. You can also perform multiple simultaneous live migrations to quickly move many virtual machines within a cluster. This provides advantages particularly when placing heavily laden hosts in maintenance mode, which causes VMs to evacuate the host. With Windows Server 2012 R2 Hyper-V, those VMs will lease the source host, and distribute among the remaining hosts, simultaneously, making full use of the available network bandwidth, and without downtime to the workload.

## Live Migration with Compression

With Windows Server 2012 R2, there have been a number of performance enhancements that help to make the process of live migration faster, more streamlined and efficient. In larger scale deployments, such as private cloud deployments or cloud hosting providers, these performance improvements can reduce overhead on the network and CPU usage in addition to reducing the amount of time for a live migration. Hyper-V administrators can configure the appropriate live migration performance options based on their environment and requirements. Live migration performance options are configured on the host settings in the Hyper-V Manager console or via the Set-VMHost Windows PowerShell cmdlet and applies to all live migrations initiated from the host.

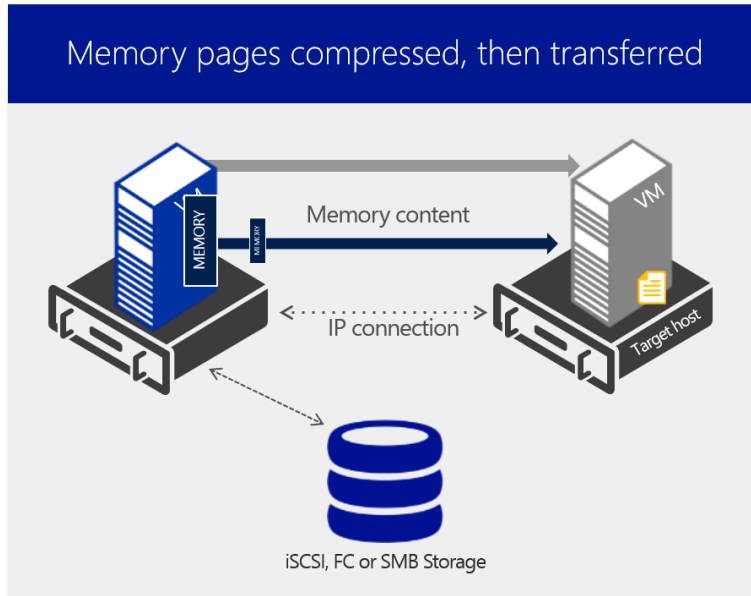


Figure 33 – Live Migration with Compression

The first of these performance enhancements, Live Migration Compression, is enabled by default, and uses spare CPU cycles on the source host to compress the memory that would need to be transferred across to the target host, as described in the process earlier. This compression process will only take place if there are spare CPU resources that it can use for the process. If a Hyper-V host is heavily burdened, with limited available CPU, compression will not be used. However, if the Hyper-V host does utilize the CPU resources to compress the memory, customers can expect to see up to a 50% reduction in time taken to live migrate a running virtual machine between hosts.

### Live Migration over SMB

Not to be confused with moving VMs whilst their virtual disks are stored on an SMB 3.0 file share, instead, Live Migration over SMB is the second of the two performance enhancements to Live Migration, and utilizes technologies from the SMB protocol to accelerate live migration to an even greater extent than that offered by compression.

With Live Migration over SMB, the memory contents of the virtual machines is copied over the network using SMB 3.0 as a transport. This means you can take advantage of some of the key SMB features to accelerate the process. **Firstly, as you add more NICs to the host, the host will utilize the SMB Multichannel capability to drive increased performance, whilst providing increased resiliency and redundancy against NIC and path failure.** **Secondly, should you invest in Remote Device Memory Access, or RDMA-capable NIC hardware (such as iWARP (10 Gbps), ROCE (10/40 Gbps), or Infiniband (56 Gbps)), SMB Direct can drive the highest levels of migration performance, on links that individually, can reach as high as 56Gbps.** Network adapters that have **RDMA**, can function at full speed with very low latency, while using very little CPU.

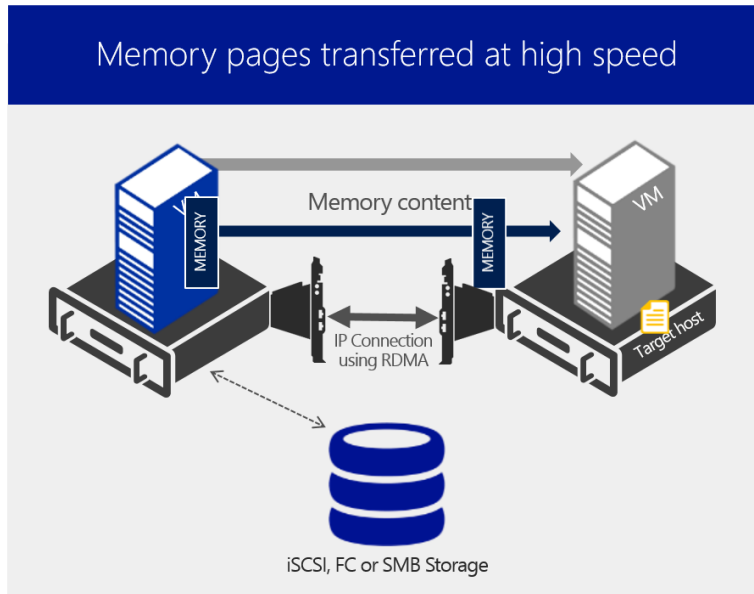


Figure 34 – Live Migration with SMB using RDMA-enabled NICs

With SMB Multichannel, SMB detects whether a network adapter has the RDMA capability, and then creates multiple RDMA connections for that single session (two per interface). This allows SMB to use the high throughput, low latency, and low CPU utilization offered by RDMA-capable network adapters. It also offers fault tolerance if you are using multiple RDMA interfaces.

With enough of these RDMA-enabled NICs in a host, all being used for Live Migration, Microsoft internal testing has seen that CPU, Disk and Network are no longer bottlenecks – in fact, the bottleneck with such fast live migration speeds, is the memory itself.

## Storage Live Migration

In Windows Server 2008 R2 Hyper-V, a virtual machine's storage could be moved only while the virtual machine was shut down. In many organizations, having the flexibility to manage storage without affecting the availability of your virtual machine workloads is a key capability. IT administrators need this flexibility to perform maintenance on storage subsystems, upgrade storage appliance firmware and software, and balance loads as capacity is used. **Windows Server 2008 R2 allowed you to move a running instance of a virtual machine by using live migration, but you still could not move the virtual machine's storage while the virtual machine was running.**

Hyper-V in Windows Server 2012 R2 provides a feature known as **Live Storage Migration**, which lets you move virtual hard disks attached to a running virtual machine. Through this feature, you can transfer virtual hard disks, with no downtime, to a new location for upgrading or migrating storage, performing backend storage maintenance, or redistributing your storage load. You can perform this operation by using a new wizard in Hyper-V Manager or the new Hyper-V cmdlets for Windows PowerShell. Live storage migration is available for both storage area network (SAN)-based and file-based storage.

When you move a running virtual machine's virtual hard disks, Hyper-V performs the following steps to move storage:

1. Throughout most of the move operation, disk reads and writes go to the source virtual hard disk.
2. While reads and writes occur on the source virtual hard disk, the disk contents are copied to the new destination virtual hard disk.

3. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.
4. After the source and destination virtual hard disks are synchronized, the virtual machine switches over to using the destination virtual hard disk.
5. The source virtual hard disk is deleted.

Just as virtual machines might need to be dynamically moved in a datacenter, allocated storage for running virtual hard disks might sometimes need to be moved for storage load distribution, storage device services, or other reasons.

Updating the physical storage that is available to Hyper-V is the most common reason for moving a virtual machine's storage. You also may want to move virtual machine storage between physical storage devices, at runtime, to take advantage of new, lower-cost storage that is supported in this version of Hyper-V (such as SMB-based storage), or to respond to reduced performance that results from bottlenecks in the storage throughput. **Windows Server 2012 R2 provides the flexibility to move virtual hard disks both on shared storage subsystems and on non-shared storage as long as a Windows Server 2012 R2 SMB 3.0 network shared folder is visible to both Hyper-V hosts.**

You can add physical storage to either a stand-alone system or to a Hyper-V cluster, and then move the virtual machine's virtual hard disks to the new physical storage while the virtual machine continues to run.

Storage migration, combined with live migration, also lets you move a virtual machine between hosts on different servers that are not using the same storage. For example, if two Hyper-V servers are each configured to use different storage devices and a virtual machine must be migrated between these two servers, you can use storage migration to a shared folder on a file server that is accessible to both servers and then migrate the virtual machine between the servers (because they both have access to that share). Following the live migration, you can use another storage migration to move the virtual hard disk to the storage that is allocated for the target server, or use "Shared Nothing" Live Migration.

### Shared Nothing Live Migration

We've discussed Live Migration, and in addition, a number of ways that Hyper-V in Windows Server 2012 R2 accelerates that process, using advanced technologies such as Live Migration Compression, and Live Migration over SMB. We've also discussed the flexibility that storing your virtual disks on SMB shares brings, along with being able to move the virtual disks of running virtual machines, without taking the VM down, using Storage Live Migration. There is however, one additional type of Live Migration, that takes the best of all of the key capabilities above, and combines them in such a way as to drive the migration flexibility to its highest, and that is **Shared Nothing Live Migration**. Shared Nothing Live Migration allows the IT administrator to move a running virtual machine, and its virtual disk(s), from one location to another, simultaneously, with no downtime. This unlocks scenarios such as VM migration from:

- Standalone Host with Local Storage to Standalone Host with Local Storage
- Standalone Host with Local Storage to Clustered Host with SAN Storage
- Clustered Host with SAN Storage to a different Cluster with alternative SAN Storage

These are just some of the flexible migration options that administrators gain through utilizing Shared Nothing live Migration.

There are a number of steps involved in the Shared Nothing Live Migration process. Firstly, when you perform a live migration of a virtual machine between two computers that do not share an infrastructure, the source server creates a connection with the destination server. This connection transfers the virtual

machine configuration data to the destination server. A skeleton virtual machine is set up on the destination server and memory is allocated to the destination virtual machine

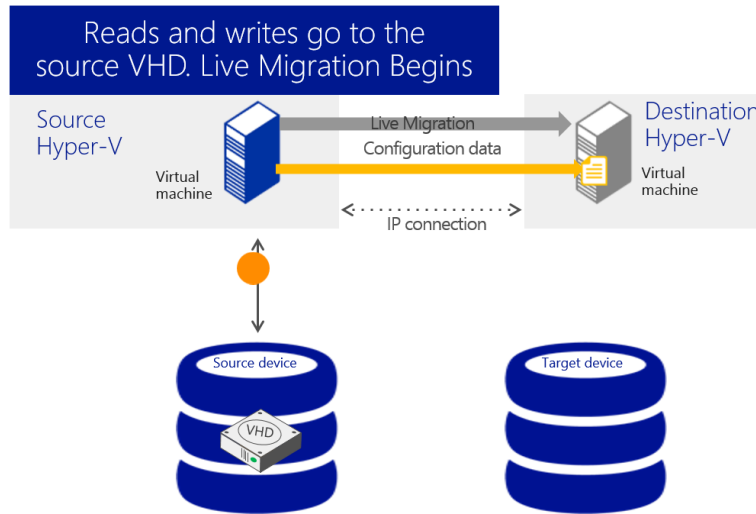


Figure 36 – Initializing a Shared Nothing Live Migration

Once the configuration data has been successfully transmitted, the disk contents are transmitted.

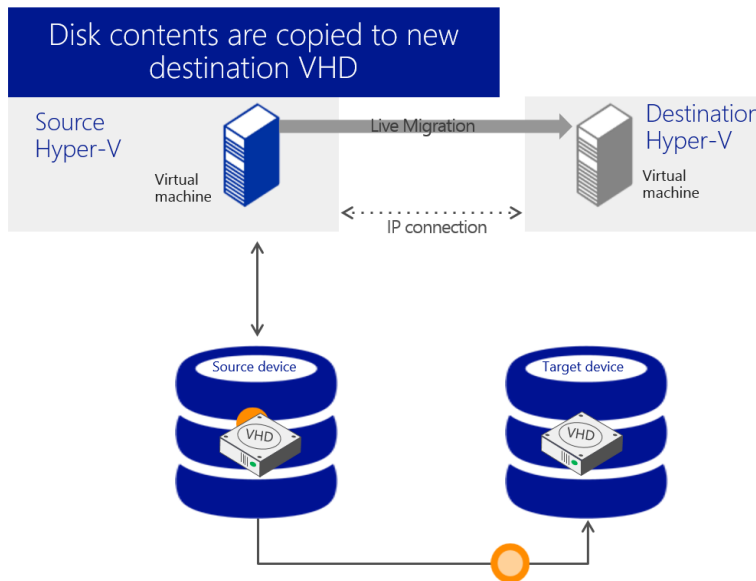


Figure 37 – Shared Nothing Live Migration – Disk contents transferred.

During this process, reads and writes are still going to the source virtual hard disk. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.

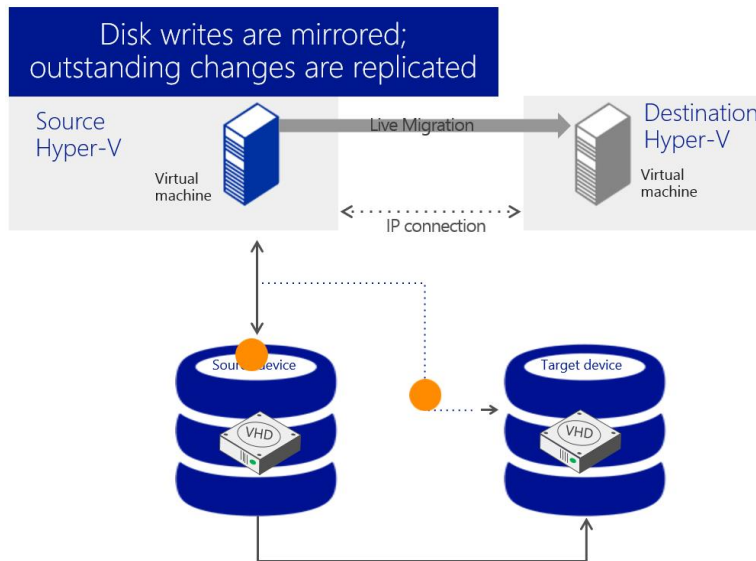


Figure 38 – Shared Nothing Live Migration – Writes are mirrored

After the source and destination virtual hard disks are completely synchronized, the virtual machine live migration is initiated, following the same process that is used for live migration with shared storage.

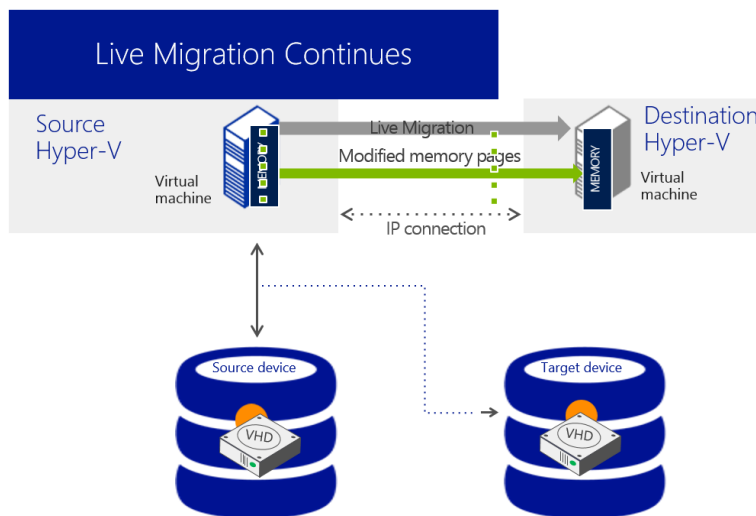


Figure 39 – Shared Nothing Live Migration – Writes are mirrored

Once the live migration is complete and the virtual machine is successfully running on the destination server, the files on the source server are deleted.



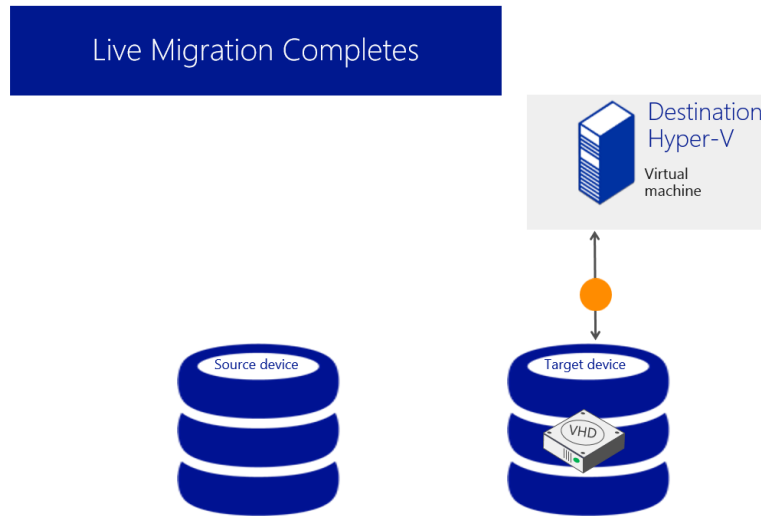


Figure 40 – Shared Nothing Live Migration – Writes are mirrored

## Cross-Version Live Migration

In previous releases of Windows Server, to move to the next version of the platform incurred downtime for key workloads, as the virtual machines were exported from old, and imported to new. With Windows Server 2012 R2 Hyper-V however, customers can upgrade from Windows Server 2012 Hyper-V with no virtual machine downtime, enabling more seamless and efficient migration of those key workloads. Note, it's important to call out that a down-level migration is not supported.

Cross-Version Live Migration allows migration between both standalone, and clustered hosts, and the process can be fully automated using PowerShell.

### Requirements

To take advantage of the discussed Live Migration capabilities, you will require the following:

- Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2
- Two or more Hyper-V hosts that support hardware virtualization, and use processors from the same manufacturer (for example, all AMD or all Intel).
- Hyper-V hosts that are part of the same Active Directory domain.
- Virtual machines configured to use virtual hard disks or virtual Fibre Channel disks
- A private network for live migration network traffic.
- Live Migration over SMB with SMB Direct requires specific **RDMA**-capable hardware.

Live migration in a cluster requires the following:

- The Windows Failover Clustering feature enabled and configured.
- CSV storage in the cluster enabled.

Live migration by using SMB shared storage requires the following:

- All files on a virtual machine (such as virtual hard disks, snapshots, and configuration) stored on a SMB 3.0 share.
- Permissions on the SMB share configured to grant access to the computer accounts of all Hyper-V hosts.

Live migration with no shared infrastructure has no additional requirements.

## Why This Matters

**Live migration, which was first introduced with Windows Server 2008 R2, was a valuable improvement for cloud management, giving organizations the ability to move virtual machines without shutting them down. As an organization's customer base grows, however, managing the cloud environment becomes more challenging because effective resource use requires administrators to move virtual machines within a cluster and between clusters more frequently.**

**With the live migration improvements in Windows Server 2012 R2 Hyper-V, organizations can now not only perform live migrations, but also move many virtual machines, and their storage, quickly and without downtime, between clusters and now even to servers that do not share storage. These improvements significantly increase the flexibility of virtual machine placement by providing truly dynamic mobility of virtual machines across a datacenter. These improvements also increase administrator efficiency and eliminate the user downtime that was previously incurred for migrations across cluster boundaries. In addition to saving time because migration speed is faster, you also save time because you can perform multiple simultaneous live migrations.**

**In addition, the ability to harness the live migration capabilities for moving workloads onto a newer, more feature rich platform, ensures that customers can embrace the new features without incurring downtime for those key workloads and applications.**

## Reliably Import Virtual Machines

Importing a virtual machine from one physical host to another can expose file incompatibilities and other unforeseen complications. Administrators often think of a virtual machine as a single, stand-alone entity that they can move to address their operational needs. In reality, a virtual machine consists of several parts:

- Virtual hard disks, stored as files in the physical storage.
- Virtual machine checkpoints, stored as a special type of virtual hard disk file.
- The saved state of the different, host-specific devices.
- The memory file, or checkpoint, for the virtual machine.
- The virtual machine configuration file, which organizes the preceding components and arranges them into a working virtual machine.

Each virtual machine, and each checkpoint associated with it, use unique identifiers. Additionally, virtual machines store and use some host-specific information, such as the path that identifies the location for virtual hard disk files. When Hyper-V starts a virtual machine, it undergoes a series of validation checks before being started. Problems such as hardware differences that might exist when a virtual machine is imported to another host can cause these validation checks to fail. That, in turn, prevents the virtual machine from starting.

Windows Server 2012 R2 includes an Import wizard that helps you quickly and reliably import virtual machines from one server to another. The Import Wizard detects and fixes more than 40 different types of incompatibilities. You don't have to worry ahead of time about the configuration that is associated with

physical hardware, such as memory, virtual switches, and virtual processors. The Import Wizard guides you through the steps to resolving incompatibilities when you import the virtual machine onto the new host.

In addition, you no longer need to export a virtual machine to be able to import it. You can simply copy a virtual machine and its associated files to the new host, and then use the Import Wizard to specify the location of the files. This “registers” the virtual machine with Hyper-V and makes it available for use. You also can recover virtual machines if the system drive fails, as long as the data drive that stores the virtual machines is intact. Automation support is also available through the Hyper-V module for Windows PowerShell, which includes cmdlets for importing virtual machines.

When you import a virtual machine, the wizard does the following:

1. **Creates a copy of the virtual machine configuration file.** This is created as a precaution in case an unexpected restart occurs on the host, such as from a power outage.
2. **Validates hardware.** Information in the virtual machine configuration file is compared to hardware on the new host.
3. **Compiles a list of errors.** This list identifies what needs to be reconfigured and determines which pages appear next in the wizard.
4. **Displays the relevant pages, one category at a time.** The wizard identifies incompatibilities to help you reconfigure the virtual machine so that it is compatible with the new host.
5. **Removes the copy of the configuration file.** After the wizard does this, the virtual machine is ready to start.

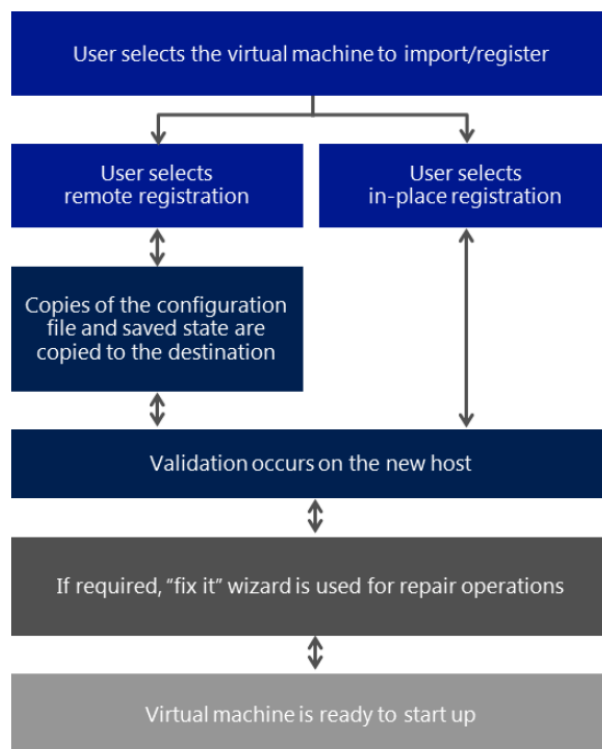


Figure 41 – VM Import Process

## Requirements

To use the Import Wizard, you need the following:

- Two installations of Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2.
- A computer that has processor support for hardware virtualization.
- A virtual machine.
- **A user account that belongs to the local Hyper-V Administrators group.**

### Why This Matters

**The new Import Wizard is a simpler, better way to import or copy virtual machines. The wizard detects and fixes potential problems, such as hardware or file differences that might exist when a virtual machine is imported to another host. As an added safety feature, the wizard creates a temporary copy of a virtual machine configuration file in case an unexpected restart occurs on the host, such as from a power outage. Windows PowerShell cmdlets for importing virtual machines let you automate the process. All of these capabilities ensure that if you need to transfer virtual machines offline, from one environment to another, the process will work seamlessly, and efficiently.**

## Automation Support for Hyper-V

PowerShell is the scripting solution for automating tasks in Windows Server. However, in earlier versions of Windows Server, writing scripts for Hyper-V with in-box tools required you to learn WMI, which provides a very flexible set of interfaces designed for developers. IT professionals who are involved with virtualization need ways to easily automate a number of administrative tasks without having to learn developer skills.

Hyper-V in Windows Server 2012 R2 provides more than 140 Hyper-V cmdlets for PowerShell to enable you to automate key tasks and processes that involve Hyper-V.

The new Hyper-V cmdlets for Windows PowerShell are intentionally designed for IT professionals and let you perform available tasks in the graphic user interface (GUI) of Hyper-V Manager and several tasks exclusively through the cmdlets in Windows PowerShell. This design is reflected in several ways.

### Task-oriented interface

Hyper-V cmdlets make it easier for IT professionals to go from thinking about the task to actually performing the task. The following table shows the task and the associated cmdlet syntax.

Task	PowerShell Command
Create a new virtual machine named "test."	New-VM -Name Test
Get a list of all virtual machines.	Get-VM
Create a new virtual hard disk at d:\VHDs\test.vhd.	New-VHD -Path D:\VHDs\test.vhd
Start all virtual machines whose name begins with "web."	Start-VM -Name web*
Connect the virtual network adapter on the "test" virtual machine to the "QA" switch.	Connect-VMNetworkAdapter -VMName test -SwitchName QA

Table 5 – Tasks and cmdlet syntax

Hyper-V administrators often must manage more than just Hyper-V. By using the same verbs as other Windows cmdlets, the Hyper-V cmdlets make it easier for administrators to extend their existing knowledge of Windows PowerShell. For example, administrators who are familiar with managing services by using Windows PowerShell can reuse the same verbs to perform the corresponding tasks on a virtual machine, as shown in the following table.

Task	Service cmdlet	Hyper-V cmdlet
Get	Get-Service	Get-VM
Configure	Set-Service	Set-VM
Create	New-Service	New-VM
Start	Start-Service	Start-VM
Stop	Stop-Service	Stop-VM
Restart	Restart-Service	Restart-VM
Suspend	Suspend-Service	Suspend-VM
Resume	Resume-Service	Resume-VM

Table 6 – Hyper-V cmdlets

There are similar examples with other core Windows PowerShell cmdlets as well, as shown in the following table.

Core PowerShell cmdlet	Hyper-V cmdlet
Import-Csv	Import-VM
Export-Csv	Export-VM
Enable-PSRemoting	Enable-VMMigration
Checkpoint-Computer	Checkpoint-VM
Measure-Command	Measure-VM

Table 7 – Windows PowerShell cmdlets

### Consistent cmdlet nouns simplify discoverability

There are many cmdlets to learn (more than 140). The nouns of the Hyper-V cmdlets make it easier for you to discover the cmdlets that they need when they need them. All cmdlets in the Hyper-V module use one of three noun prefixes in the following table.

Prefi x	Purpose
VM	Cmdlets for managing virtual machines
VHD	Cmdlets for managing virtual hard disk files
VFD	Cmdlets for managing virtual floppy disk files

## Requirements

To use the new Hyper-V cmdlets in Windows Server 2012 R2, you'll need the following:

- Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2
- Administrator or Hyper-V Administrator user account.

Optionally, if you want to use the Hyper-V cmdlets remotely, you can install the Hyper-V Windows PowerShell cmdlets feature on a computer running Windows 8.1, and run the cmdlets as an Administrator or Hyper-V Administrator on the server.

### Why This Matters

**Before Windows Server 2012 and subsequently, Windows Server 2012 R2, automation of Hyper-V management tasks required writing scripts using WMI, a skill that many datacenter administrators do not have, thus making the automation difficult. Windows Server 2012 R2 provides a rich, powerful, comprehensive, and simple-to-learn set of Windows PowerShell cmdlets, datacenter administrators have an easier time using cmdlets to automate most Hyper-V tasks (such as creating a new virtual machine, importing and exporting virtual machines, and connecting a virtual network adaptor to a virtual machine). You can use these new cmdlets to automate basic and complex datacenter tasks with ease and reduce the administrative overhead in your cloud computing environment.**

## Hyper-V Network Virtualization

Hyper-V Network Virtualization provides “virtual networks” (called a VM network) to virtual machines similar to how server virtualization (hypervisor) provides “virtual machines” to the operating system. Network virtualization decouples virtual networks from the physical network infrastructure and removes the constraints of VLAN and hierarchical IP address assignment from virtual machine provisioning. This flexibility makes it easy for customers to move to IaaS clouds and efficient for hosters and datacenter administrators to manage their infrastructure, while maintaining the necessary multi-tenant isolation, security requirements, and supporting overlapping Virtual Machine IP addresses.




Figure 42 – Server virtualization versus network virtualization

Customers want to seamlessly extend their datacenters to the cloud. Today there are technical challenges in making such seamless hybrid cloud architectures. One of the biggest hurdles customers face is reusing their existing network topologies (subnets, IP addresses, network services, and so on.) in the cloud and bridging between their on-premise resources and their cloud resources. Hyper-V Network Virtualization provides the concept of a VM Network that is independent of the underlying physical network. With this concept of a VM Network, composed of one or more Virtual Subnets, the exact location in the physical network of virtual machines attached to a virtual network is decoupled from the virtual network topology. As a result, customers can easily move their virtual subnets to the cloud while preserving their existing IP addresses and topology in the cloud so that existing services continue to work unaware of the physical location of the subnets. That is, Hyper-V Network Virtualization enables a seamless hybrid cloud.

In addition to hybrid cloud, many organizations are consolidating their datacenters and creating private clouds to internally get the efficiency and scalability benefit of cloud architectures. Hyper-V Network Virtualization allows better flexibility and efficiency for private clouds by decoupling a business unit's network topology (by making it virtual) from the actual physical network topology. In this way, the business units can easily share an internal private cloud while being isolated from each other and continue to keep existing network topologies. The datacenter operations team has flexibility to deploy and dynamically move workloads anywhere in the datacenter without server interruptions providing better operational efficiencies and an overall more effective datacenter.

For workload owners, the key benefit is that they can now move their workload "topologies" to the cloud without changing their IP addresses or re-writing their applications. For example, the typical three-tier LOB application is composed of a front end tier, a business logic tier, and a database tier. Through policy, Hyper-V Network Virtualization allows customer onboarding all or parts of the three tiers to the cloud, while keeping the routing topology and the IP addresses of the services (i.e. virtual machine IP addresses), without requiring the applications to be changed.

For infrastructure owners, the additional flexibility in virtual machine placement makes it possible to move workloads anywhere in the datacenters without changing the virtual machines or reconfiguring the networks. For example Hyper-V Network Virtualization enables cross subnet live migration so that a virtual machine can live migrate anywhere in the datacenter without a service disruption. Previously live migration was limited to the same subnet restricting where virtual machines could be located. Cross subnet live migration allows administrators to consolidate workloads based on dynamic resource requirements, energy

efficiency, and can also accommodate infrastructure maintenance without disrupting customer workload up time.

## Practical Applications

With the success of virtualized datacenters, IT organizations and hosting providers (providers who offer colocation or physical server rentals) have begun offering more flexible virtualized infrastructures that make it easier to offer on-demand server instances to their customers. This new class of service is referred to as Infrastructure as a Service (IaaS). Windows Server 2012 R2 provides all the required platform capabilities to enable enterprise customers to build private clouds and transition to an IT as a service operational model. Windows Server 2012 R2 also enables hosters to build public clouds and offer IaaS solutions to their customers. When combined with System Center Virtual Machine Manager to manage Hyper-V Network Virtualization policy, Microsoft provides a powerful cloud solution.

Windows Server 2012 R2 Hyper-V Network Virtualization provides policy-based, software-controlled network virtualization that reduces the management overhead faced by enterprises when they expand dedicated IaaS clouds, and it provides cloud hosters better flexibility and scalability for managing virtual machines to achieve higher resource utilization.

An IaaS scenario that has virtual machines from different organizational divisions (dedicated cloud) or different customers (hosted cloud) requires secure isolation. Today's solution, virtual local area networks (VLANs), can present significant disadvantages in this scenario.

**VLANs** - Currently, VLANs are the mechanism that most organizations use to support address space reuse and tenant isolation. A VLAN uses explicit tagging (VLAN ID) in the Ethernet frame headers, and it relies on Ethernet switches to enforce isolation and restrict traffic to network nodes with the same VLAN ID. The main disadvantages with VLANs are as follows:

- Increased risk of an inadvertent outage due to cumbersome reconfiguration of production switches whenever virtual machines or isolation boundaries move in the dynamic datacenter.
- Limited in scalability because there is a maximum of 4094 VLANs and typical switches support no more than 1000 VLAN IDs.
- Constrained within a single IP subnet, which limits the number of nodes within a single VLAN and restricts the placement of virtual machines based on physical locations. Even though VLANs can be expanded across sites, the entire VLAN must be on the same subnet.

**IP address assignment** - In addition to the disadvantages that are presented by VLANs, virtual machine IP address assignment presents issues, which include:

- Physical locations in datacenter network infrastructure determine virtual machine IP addresses. As a result, moving to the cloud typically requires changing IP addresses of the service workloads.
- Policies are tied to IP addresses, such as firewall rules, resource discovery and directory services, and so on. Changing IP addresses requires updating all the associated policies.
- Virtual machine deployment and traffic isolation are dependent on the topology.

When datacenter network administrators plan the physical layout of the datacenter, they must make decisions about where subnets will be physically placed and routed. These decisions are based on IP and Ethernet technology that influence the potential IP addresses that are allowed for virtual machines running on a given server or a blade that is connected to a particular rack in the datacenter. When a virtual machine is provisioned and placed in the datacenter, it must adhere to these choices and restrictions regarding the IP address. Therefore, the typical result is that the datacenter administrators assign new IP addresses to the virtual machines.



The problem with this requirement is that in addition to being an address, there is semantic information associated with an IP address. For instance, one subnet may contain given services or be in a distinct physical location. Firewall rules, access control policies, and IPsec security associations are commonly associated with IP addresses. Changing IP addresses forces the virtual machine owners to adjust all their policies that were based on the original IP address. This renumbering overhead is so high that many enterprises choose to deploy only new services to the cloud, leaving legacy applications alone.

Hyper-V Network Virtualization decouples virtual networks for customer virtual machines from the physical network infrastructure. As a result, it enables customer virtual machines to maintain their original IP addresses, while allowing datacenter administrators to provision customer virtual machines anywhere in the datacenter without reconfiguring physical IP addresses or VLAN IDs.

## Network Virtualization - Key Benefits

- **Enables flexible workload placement – Network isolation and IP address re-use without VLANs** - Hyper-V Network Virtualization decouples the customer's virtual networks from the physical network infrastructure of the hosters, providing freedom for workload placements inside the datacenters. Virtual machine workload placement is no longer limited by the IP address assignment or VLAN isolation requirements of the physical network because it is enforced within Hyper-V hosts based on software-defined, multitenant virtualization policies.

Virtual machines from different customers with overlapping IP addresses can now be deployed on the same host server without requiring cumbersome VLAN configuration or violating the IP address hierarchy. This can streamline the migration of customer workloads into shared IaaS hosting providers, allowing customers to move those workloads without modification, which includes leaving the virtual machine IP addresses unchanged. For the hosting provider, supporting numerous customers who want to extend their existing network address space to the shared IaaS datacenter is a complex exercise of configuring and maintaining isolated VLANs for each customer to ensure the coexistence of potentially overlapping address spaces. With Hyper-V Network Virtualization, supporting overlapping addresses is made easier and requires less network reconfiguration by the hosting provider.

In addition, physical infrastructure maintenance and upgrades can be done without causing a down time of customer workloads. With Hyper-V Network Virtualization, virtual machines on a specific host, rack, subnet, VLAN, or entire cluster can be migrated without requiring a physical IP address change or major reconfiguration.

- **Enables easier moves for workloads to a shared IaaS cloud** - With Hyper-V Network Virtualization, IP addresses and virtual machine configurations remain unchanged. This enables IT organizations to more easily move workloads from their datacenters to a shared IaaS hosting provider with minimal reconfiguration of the workload or their infrastructure tools and policies. In cases where there is connectivity between two datacenters, IT administrators can continue to use their tools without reconfiguring them.
- **Enables live migration across subnets** - Live migration of virtual machine workloads traditionally has been limited to the same IP subnet or VLAN because crossing subnets required the virtual machine's guest operating system to change its IP address. This address change breaks existing communication and disrupts the services running on the virtual machine. With Hyper-V Network Virtualization, workloads can be live migrated from servers running Windows Server 2012 in one subnet to servers running Windows Server 2012 in a different subnet without changing the workload IP addresses. Hyper-V Network Virtualization ensures that virtual machine location changes due to live migration are updated and synchronized among hosts that have ongoing communication with the migrated virtual machine.

- **Enables easier management of decoupled server and network administration** - Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, and network administrators can focus on overall network infrastructure and traffic management. This enables datacenter server administrators to deploy and migrate virtual machines without changing the IP addresses of the virtual machines. There is reduced overhead because Hyper-V Network Virtualization allows virtual machine placement to occur independently of network topology, reducing the need for network administrators to be involved with placements that might change the isolation boundaries.
- **Simplifies the network and improves server/network resource utilization** - The rigidity of VLANs and the dependency of virtual machine placement on a physical network infrastructure results in overprovisioning and underutilization. By breaking the dependency, the increased flexibility of virtual machine workload placement can simplify the network management and improve server and network resource utilization. Note that Hyper-V Network Virtualization supports VLANs in the context of the physical datacenter. For example, a datacenter may want all Hyper-V Network Virtualization traffic to be on a specific VLAN.
- **Is compatible with existing infrastructure and emerging technology** - Hyper-V Network Virtualization can be deployed in today's datacenter, yet it is compatible with emerging datacenter "flat network" technologies.
- **Provides for interoperability and ecosystem readiness** - Hyper-V Network Virtualization supports multiple configurations for communication with existing resources, such as cross premise connectivity, storage area network (SAN), non-virtualized resource access, and so on. Microsoft is committed to working with ecosystem partners to support and enhance the experience of Hyper-V Network Virtualization in terms of performance, scalability, and manageability.
- **Uses Windows PowerShell and WMI** - Hyper-V Network Virtualization supports Windows PowerShell and Windows Management Instrumentation (WMI) for configuring the network virtualization and isolation policies. The Windows PowerShell cmdlets for Hyper-V Network Virtualization enable administrators to build command-line tools or automated scripts to configure, monitor, and troubleshoot network isolation policies.

## Network Virtualization Concepts

In Hyper-V Network Virtualization (HNV), a customer is defined as the "owner" of a group of virtual machines that are deployed in a datacenter. A customer can be a corporation or enterprise in a multitenant public datacenter, or a division or business unit within a private datacenter. Each customer can have one or more VM networks in the datacenter, and each VM network consists of one or more virtual subnets.

### VM network

- Each VM network consists of one or more virtual subnets. A VM network forms an isolation boundary where the virtual machines within a VM network can communicate with each other. As a result, virtual subnets in the same VM network must not use overlapping IP address prefixes.
- Each VM network has a Routing Domain which identifies the VM network. The Routing Domain ID (RDID), which identifies the VM network, is assigned by datacenter administrators or datacenter management software, such as System Center 2012 R2 Virtual Machine Manager (VMM). The RDID is a Windows GUID — for example, "{11111111-2222-3333-4444-000000000000}".

### Virtual subnets

- A virtual subnet implements the Layer 3 IP subnet semantics for the virtual machines in the same virtual subnet. The virtual subnet is a broadcast domain (similar to a VLAN). Virtual machines in the same virtual subnet must use the same IP prefix.
- Each virtual subnet belongs to a single VM network (RDID), and it is assigned a unique Virtual Subnet ID (VSID). The VSID must be unique within the datacenter and is in the range 4096 to  $2^{24}-2$ .

A key advantage of the VM network and routing domain is that it allows customers to bring their network topologies to the cloud. The figure below shows an example where the Contoso Corp has two separate networks, the R&D Net and the Sales Net. Because these networks have different routing domain IDs, they cannot interact with each other. That is, Contoso R&D Net is isolated from Contoso Sales Net even though both are owned by Contoso Corp. Contoso R&D Net contains three virtual subnets. Note that both the RDID and VSID are unique within a datacenter

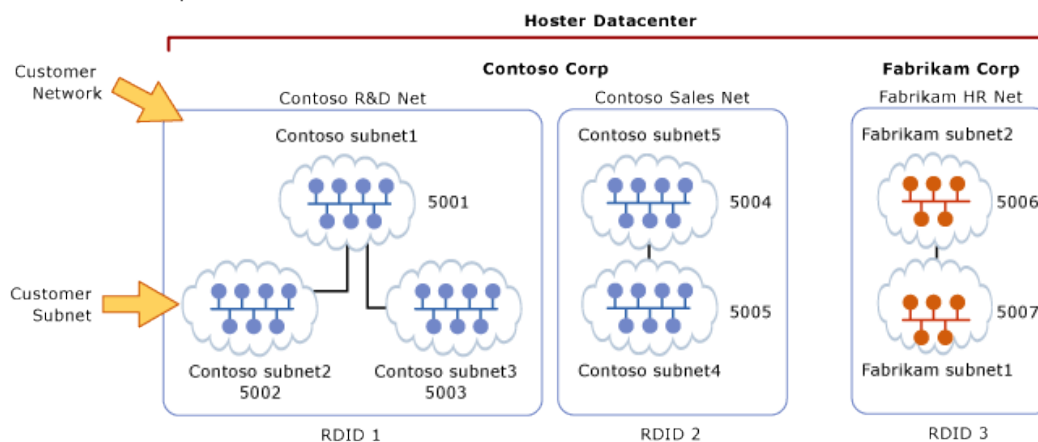


Figure 43 – Customer networks and virtual subnets

In the above figure, the virtual machines with VSID 5001 can have their packets routed or forwarded by HNV to virtual machines with VSID 5002 or VSID 5003. Before delivering the packet to the Hyper-V switch, HNV will update the VSID of the incoming packet to the VSID of the destination virtual machine. This will only happen if both VSIDs are in the same RDID. If the VSID that is associated with the packet does not match the VSID of the destination virtual machine, the packet will be dropped. Therefore, virtual network adapters with RDID1 cannot send packets to virtual network adapters with RDID2.

**Note** - In the packet flow description above, the term “virtual machine” actually means the “virtual network adapter” on the virtual machine. The common case is that a virtual machine only has a single virtual network adapter. In this case, the words virtual machine and virtual network adapter can conceptually mean the same thing. Because a virtual machine can have multiple virtual network adapters, and these virtual network adapters can have different VirtualSubnetIDs (VSIDs) or RoutingDomainIDs (RDIDs), HNV specifically focuses on the packets sent and received between virtual network adapters.

Each virtual subnet defines a Layer 3 IP subnet and a Layer 2 (L2) broadcast domain boundary similar to a VLAN. When a virtual machine broadcasts a packet, this broadcast is limited to the virtual machines that are attached to switch ports with the same VSID. Each VSID can be associated with a multicast address in the PA. All broadcast traffic for a VSID is sent on this multicast address.

**Note** - HNV does NOT depend on broadcast or multicast. For broadcast or multicast packets in a VM network, a PA multicast IP address is used if configured. However, the many datacenter operators do not enable multicast in their environments. As a result, when a PA multicast address is not available an intelligent PA unicast replication is used. This means that packets are unicasted only to PA addresses that are

configured for the particular virtual subnet the packet is on. In addition, only one unicast packet per host is sent no matter how many relevant virtual machines are on the host.

In addition to being a broadcast domain, the VSID provides isolation. A virtual network adapter in HNV is connected to a Hyper-V switch port that has a VSID ACL. If a packet arrives on this Hyper-V switch port with a different VSID the packet is dropped. Packets will only be delivered on a Hyper-V switch port if the VSID of the packet matches the VSID of the switch port. This is the reason, in the above example, that packets flowing from VSID 5001 to 5003 must have the VSID in the packet modified before delivery to the destination virtual machine.

If the Hyper-V switch port does not have a VSID ACL, the virtual network adapter that is attached to that switch port is not part of a HNV virtual subnet. Packets sent from a virtual network adapter that does not have a VSID ACL will pass unmodified through the Hyper-V switch.

When a virtual machine sends a packet, the VSID of the Hyper-V switch port is associated with this packet. On the receiving side, HNV delivers to the Hyper-V switch the VSID in the OOB along with the decapsulated packet. On the receiving end, HNV performs a policy lookup and adds the VSID to the OOB data before the packet is passed to the Hyper-V switch.

**Note** - Hyper-V Switch Extensions can operate in both the Provider Address (PA) space and the Customer Address (CA) space. This means the VSID is available to the switch extensions. This allows the switch extension to become multitenant aware. For example, a firewall switch extension can differentiate CA IP address 10.1.1.5 with OOB containing VSID 5001 from the same CA IP address with VSID 6001.

## Packet Encapsulation

Each virtual network adapter in HNV is associated with two IP addresses:

- **Customer Address (CA)** - The IP address that is assigned by the customer, based on their intranet infrastructure. This address enables the customer to exchange network traffic with the virtual machine as if it had not been moved to a public or private cloud. The CA is visible to the virtual machine and reachable by the customer.
- **Provider Address (PA)** - The IP address that is assigned by the hoster or the datacenter administrators based on their physical network infrastructure. The PA appears in the packets on the network that are exchanged with the server running Hyper-V that is hosting the virtual machine. The PA is visible on the physical network, but not to the virtual machine.

The CAs maintain the customer's network topology, which is virtualized and decoupled from the actual underlying physical network topology and addresses, as implemented by the PAs. The following diagram shows the conceptual relationship between virtual machine CAs and network infrastructure PAs as a result of network virtualization.

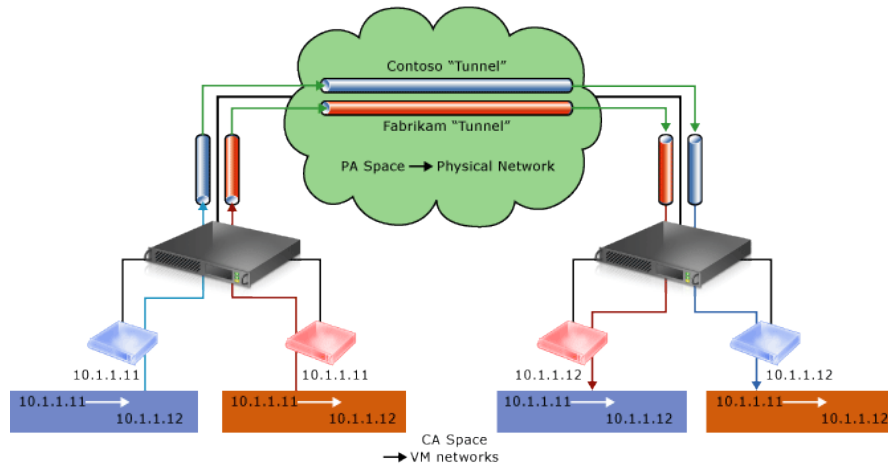


Figure 44 – Conceptual diagram of network virtualization over physical infrastructure

In the diagram, customer virtual machines are sending data packets in the CA space, which traverse the physical network infrastructure through their own virtual networks, or “tunnels”. In the example above, the tunnels can be thought of as “envelopes” around the Contoso and Fabrikam data packets with green shipping labels (PA addresses) to be delivered from the source host on the left to the destination host on the right. The key is how the hosts determine the “shipping addresses” (PA’s) corresponding to the Contoso and the Fabrikam CA’s, how the “envelope” is put around the packets, and how the destination hosts can unwrap the packets and deliver to the Contoso and Fabrikam destination virtual machines correctly.

This simple analogy highlighted the key aspects of network virtualization:

- Each virtual machine CA is mapped to a physical host PA. There can be multiple CAs associated with the same PA.
- Virtual machines send data packets in the CA spaces, which are put into an “envelope” with a PA source and destination pair based on the mapping.
- The CA-PA mappings must allow the hosts to differentiate packets for different customer virtual machines.

As a result, the mechanism to virtualize the network is to virtualize the network addresses used by the virtual machines.

## Network Virtualization through Address Virtualization

HNV supports Network Virtualization for Generic Routing Encapsulation (NVGRE) as the mechanism to virtualize the IP Address:

**Generic Routing Encapsulation** - This network virtualization mechanism uses the Generic Routing Encapsulation (NVGRE) as part of the tunnel header. In NVGRE, the virtual machine’s packet is encapsulated inside another packet. The header of this new packet has the appropriate source and destination PA IP addresses in addition to the Virtual Subnet ID, which is stored in the Key field of the GRE header, as shown in the figure below.

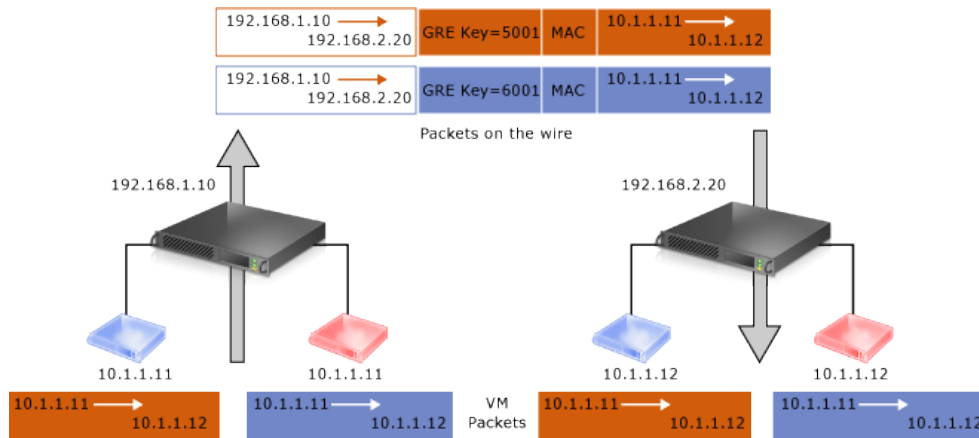


Figure 45 – Network virtualization - NVGRE encapsulation

The Virtual Subnet ID allows hosts to identify the customer virtual machine for any given packet, even though the PA's and the CA's on the packets may overlap. This allows all virtual machines on the same host to share a single PA, as shown in the figure above.

Sharing the PA has a big impact on network scalability. The number of IP and MAC addresses that need to be learned by the network infrastructure can be substantially reduced. For instance, if every end host has an average of 30 virtual machines, the number of IP and MAC addresses that need to be learned by the networking infrastructure is reduced by a factor of 30. The embedded Virtual Subnet IDs in the packets also enable easy correlation of packets to the actual customers.

With Windows Server 2012 and Windows Server 2012 R2, HNV fully supports NVGRE out of the box; it does NOT require upgrading or purchasing new network hardware such as NICs (Network Adapters), switches, or routers. This is because the NVGRE packet on the wire is a regular IP packet in the PA space, which is compatible with today's network infrastructure.

Windows Server 2012 made working with standards a high priority. Along with key industry partners (Arista, Broadcom, Dell, Emulex, Hewlett Packard, and Intel) Microsoft published a draft RFC that describes the use of Generic Routing Encapsulation (GRE), which is an existing IETF standard, as an encapsulation protocol for network virtualization. For more information, see the following Internet Draft: [Network Virtualization using Generic Routing Encapsulation](#). As NVGRE-aware becomes commercially available the benefits of NVGRE will become even greater.

### NVGRE Encapsulated Task Offload

High-speed network adapters implement a number of offloads (ex. Large Send Offload (LSO), Receive Side Scaling (RSS) and Virtual Machine Queue (VMQ)) that allow full utilization of the network adapter's throughput. As an example, a computer with a network adapter capable of 10 Gbps throughput might only be able to perform at 4 or 5 Gbps throughput without particular offloads enabled. In addition, even if it is capable of full throughput, the CPU utilization to perform at maximum throughput will be much higher than with offloads enabled.

For non-virtualized traffic, offloads just work. NVGRE, on the other hand, is an encapsulation protocol, which means that the network adapter must access the CA packet to perform the offload. In Windows Server 2012 R2 Hyper-V, NVGRE is the only way to virtualize traffic (in Windows Server 2012 IP Rewrite was also supported but not recommended; IP Rewrite has been removed from Windows Server 2012 R2) so NVGRE task offload becomes more important.

Microsoft has worked closely with network adaptor partners on a solution to these performance challenges, called NVGRE Encapsulated Task Offload. When a network adaptor supports NVGRE Encapsulated Task Offload it ensures that all relevant offloads work with HNV.

At TechEd 2013, two partners announced their next generation network adaptors will support NVGRE Encapsulated Task Offload. You can read the press releases from [Mellanox](#) and [Emulex](#) for more details.

## Network Virtualization Architecture

The figure below shows the architectural differences between HNV in Windows Server 2012 and in Windows Server 2012 R2. The basic change was that the HNV filter moved from being an NDIS lightweight filter (LWF) to being part of the Hyper-V virtual switch.

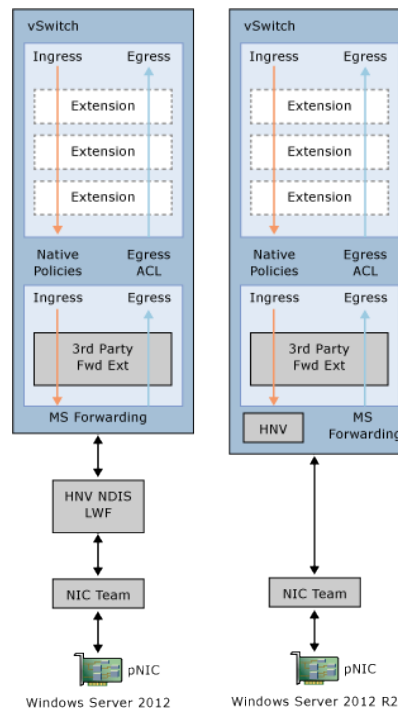


Figure 46 – Network virtualization - NVGRE encapsulation

In Windows Server 2012 HNV being an NDIS LWF meant that Hyper-V Switch extensions only worked on the customer address space. For capture and filter extensions this meant they were not aware of the underlying physical networking being used for HNV packets. For forwarding switch extensions, HNV being an NDIS LWF meant that they could not co-exist with HNV, so customer had to choose one using HNV or a particular forwarding extension. In R2, administrators can now use switch extensions on both the original customer address packet and the encapsulated provider address packet. In addition, forwarding switch extensions can co-exist with HNV allowing multiple network virtualization solutions (one provided by HNV and another provided by the forwarding switch extension) to co-exist on the same Hyper-V host.

Improved interoperability with switch extensions was the primary reason for the change but a useful side effect is that the HNVDIS LWF does not have to be bound to network adaptors anymore. Once you attach a network adaptor to the virtual switch you can enable HNV simply by assigning a Virtual Subnet ID to a particular virtual network adaptor. For those using SCVMM to manage your VM networks this is transparent but anyone using PowerShell this will save an often-missed step.

Each virtual machine network adapter is configured with an IPv4, and/or an IPv6 address. These are the CAs that will be used by the virtual machines to communicate with each other, and they are carried in the IP

packets from the virtual machines. HNV virtualizes the CAs to PAs based on the network virtualization policies.

A virtual machine sends a packet with source address CA1, which is virtualized based on HNV policy in the Hyper-V switch. A special network virtualization access control list based on VSID isolates the virtual machine from other virtual machines that are not part of the same virtual subnet or part of the same routing domain.

## Network Virtualization Routing | Between Virtual Subnets

As in physical networks, routing is an important part of HNV. There are two key aspects to understand: how packets are routed between virtual subnets and how packets are routed outside a virtual network.

In a physical network, a subnet is the Layer 2 (L2) domain where computers (virtual and physical) can directly communicate with each other without having to be routed. In Windows, if you statically configure a network adapter you can set a "default gateway" which is the IP address to send all traffic that is going out of the particular subnet so that it can be routed appropriately. This is typically the router for your physical network. HNV uses a built in router that is part of every host to form a distributed router for a virtual network. This means that every host, in particular the Hyper-V Virtual Switch, acts as the default gateway for all traffic that is going between Virtual Subnets that are part of the same VM network. In Windows Server 2012 and Windows Server 2012 R2 the address used as the default gateway is the lowest entry for the subnet (as an example, it is the ".1" address for a /24 subnet prefix). This address is reserved in each virtual subnet for the default gateway and cannot be used by virtual machines in the virtual subnet.

HNV acting as a distributed router allows for a very efficient way for all traffic inside a VM Network to be routed appropriately because each host can directly route the traffic to the appropriate host without needing an intermediary. This is particularly true when two virtual machines in the same VM Network but different Virtual Subnets are on the same physical host. As you will see later in this section, the packet never has to leave the physical host.

## Network Virtualization Routing | Outside the Virtual Network

Most customer deployments will require communication from the HNV environment to resources that are not part of the HNV environment. Network Virtualization gateways are required to allow communication between the two environments. Scenarios requiring a HNV Gateway include Private Cloud and Hybrid Cloud. Basically, HNV gateways are required for VPNs and routing.

Gateways can come in different physical form factors. They can be built on Windows Server 2012 R2, incorporated into a Top of Rack (TOR) switch, a load balancer, put into other existing network appliances, or can be a new stand-alone network appliance.

The **Windows Server Gateway (WSG)**, based on Windows Server 2012 R2, is a virtual machine-based software router that allows Cloud Service Providers (CSPs) and Enterprises to enable datacenter and cloud network traffic routing between virtual and physical networks, including the Internet.

In Windows Server 2012 R2, the WSG routes network traffic between the physical network and VM network resources, regardless of where the resources are located. You can use the WSG to route network traffic between physical and virtual networks at the same physical location or at many different physical locations. For example, if you have both a physical network and a virtual network at the same physical location, you can deploy a computer running Hyper-V that is configured with a WSG VM to route traffic between the virtual and physical networks. In another example, if your virtual networks exist in the cloud, your CSP can deploy a WSG so that you can create a virtual private network (VPN) connection between your VPN server



and the CSP's WSG; when this link is established you can connect to your virtual resources in the cloud over the VPN connection.

### Windows Server Gateway Integration with Network Virtualization

WSG is integrated with Hyper-V Network Virtualization, and is able to route network traffic effectively in circumstances where there are many different customers – or tenants – who have isolated virtual networks in the same datacenter.

Multi-tenancy is the ability of a cloud infrastructure to support the virtual machine workloads of multiple tenants, but isolate them from each other, while all of the workloads run on the same infrastructure. The multiple workloads of an individual tenant can interconnect and be managed remotely, but these systems do not interconnect with the workloads of other tenants, nor can other tenants remotely manage them.

For example, an Enterprise might have many different virtual subnets, each of which is dedicated to servicing a specific department, such as Research and Development or Accounting. In another example, a CSP has many tenants with isolated virtual subnets existing in the same physical datacenter. In both cases, WSG can route traffic to and from each tenant while maintaining the designed isolation of each tenant. This capability makes the WSG multitenant-aware.

### Clustering the Windows Server Gateway for HA

WSG is deployed on a dedicated computer that is running Hyper-V and that is configured with one VM. The VM is then configured as a WSG.

For high availability of network resources, you can deploy WSG with failover by using two physical host servers running Hyper-V that are each also running a virtual machine (VM) that is configured as a gateway. The gateway VMs are then configured as a cluster to provide failover protection against network outages and hardware failure.

When you deploy WSG, the host servers running Hyper-V and the VMs that you configure as gateways must be running Windows Server 2012 R2.

Unless otherwise noted in the illustrations that are provided in the sections below, the following icon represents two Hyper-V hosts, each of which is running a VM configured as a WSG. In addition, both the servers running Hyper-V and the VMs on each server are running Windows Server 2012 R2, and the gateway VMs are clustered.



### Private Cloud Environments

Private cloud is a computing model that uses infrastructure dedicated to your organization. A private cloud shares many of the characteristics of public cloud computing including resource pooling, self-service, elasticity, and metered services delivered in a standardized manner with the additional control and customization available from dedicated resources.

The only fundamental difference between a private cloud and a public cloud is that a public cloud provides cloud resources to multiple organizations, while the private cloud hosts resources for a single organization. However, a single organization may have multiple business units and divisions which can lend itself to being multi-tenant in nature. In these circumstances, private cloud shares many of the security and isolation requirements of public cloud.

For Enterprises that deploy an on-premises private cloud, WSG can route traffic between virtual networks and the physical network. For example, if you have created virtual networks for one or more of your departments, such as Research and Development or Accounting, but many of your key resources (such as Active Directory Domain Services, SharePoint, or DNS) are on your physical network, WSG can route traffic between the virtual network and the physical network to provide employees working on the virtual network with all of the services they need.

In the illustration below, the physical and virtual networks are at the same physical location. WSG is used to route traffic between the physical network and virtual networks.




Figure 47 – Single location, with the WSG being used to connect VMs with physical infrastructure

### Hybrid Cloud Environments

For CSPs that host many tenants in their datacenter, WSG provides a multitenant gateway solution that allows your tenants to access and manage their resources from remote sites, and that allows network traffic flow between virtual resources in your datacenter and their physical network.

In the illustration below, a CSP provides datacenter network access to multiple tenants, some of whom have multiple sites across the Internet. In this example, tenants use third party VPN servers at their corporate sites, while the CSP uses WSG for the site-to-site VPN connections.

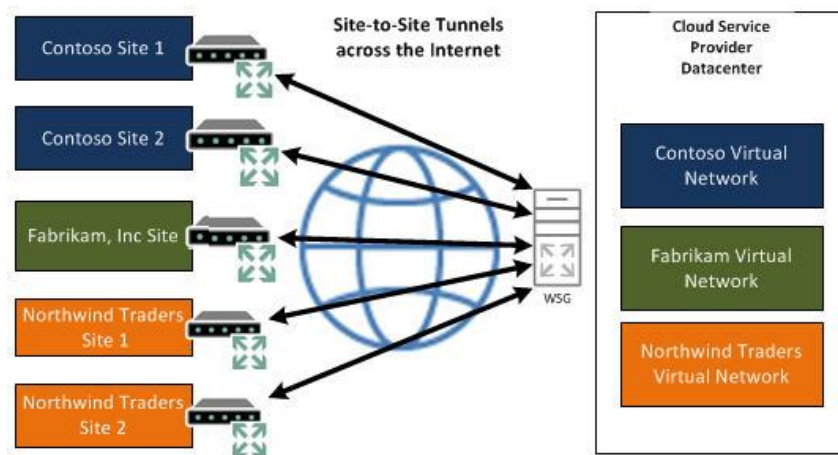


Figure 48 – CSP with using the WSG for connecting multiple customers to their hosted VM networks

These are just two of the ways that the WSG can be used to connect the outside world, to VMs running inside networks that have been created using Hyper-V Network Virtualization.

### Requirements

To fully utilize Hyper-V Network Virtualization, you will require the following:

- Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2
- System Center 2012 R2 Virtual Machine Manager
- A Windows Server 2012 R2 Hyper-V Host and Windows Server 2012 R2 VM for use as a Gateway or a Partner solution providing the Gateway functionality.

### Why This Matters

**Cloud-based datacenters can provide many benefits such as improved scalability and better resource utilization. To realize these potential benefits requires a technology that fundamentally addresses the issues of multi-tenant scalability in a dynamic environment. HNV was designed to address these issues and also improve the operational efficiency of the datacenter by decoupling the virtual network topology for the physical network topology. Building on an existing standard, HNV runs in today's datacenter and as NVGRE-aware hardware becomes available the benefits will continue to increase. Customers, with HNV can now consolidate their datacenters into a private cloud or seamlessly extend their datacenters to a hoster's environment with a hybrid cloud.**

**With the inclusion of the Windows Server Gateway, Customers can now seamlessly integrate internal address spaces with external networks, including those provided by hosters, without losing the ability to route NVGRE packets. Hosters can support multi-tenant NVGRE environments without having to run a separate VPN appliance and NAT environment for each customer. This is extremely useful in multi-tenant environments. The solution is highly available using guest clustering. Customers can enable this in-box network virtualization gateway without the need for specialized third party hardware or software, or can purchase a hardware-based solution or appliance of their choice through a vendor-partner solution**

## High Availability & Resiliency

We've spent a considerable amount of time discussing some of the key features of the platform that provide for immense scalability and performance, security, and most recently, features that enable complete flexibility from both a VM migration perspective, but also a networking perspective. One thing we do have to account for however, is what happens when things go wrong. What happens when a piece of hardware, such as a NIC fails? What about a host, or even an entire datacenter? Fortunately, Windows Server 2012 R2 has a number of key features and capabilities that provide resiliency at each of those different levels, ensuring you can virtualize your mission critical, high performance workloads, and be confident that they are providing a high level of continuous service to the business.

### NIC Teaming

The failure of an individual Hyper-V port or virtual network adapter can cause a loss of connectivity for a virtual machine. Using multiple virtual network adapters in a Network Interface Card (NIC) Teaming solution can prevent connectivity loss and, when multiple adapters are connected, multiply throughput.

To increase reliability and performance in virtualized environments, Windows Server 2012 R2 includes built-in support for NIC Teaming-capable network adapter hardware. Although NIC Teaming in Windows Server 2012 R2 is not a Hyper-V feature, it is important for business-critical Hyper-V environments because it can provide increased reliability and performance for virtual machines. NIC Teaming is also known as “network adapter teaming technology” and “load balancing failover” (LBFO).

## Architecture

Today, all NIC Teaming solutions on the market have a similar architecture, as shown in the figure below.

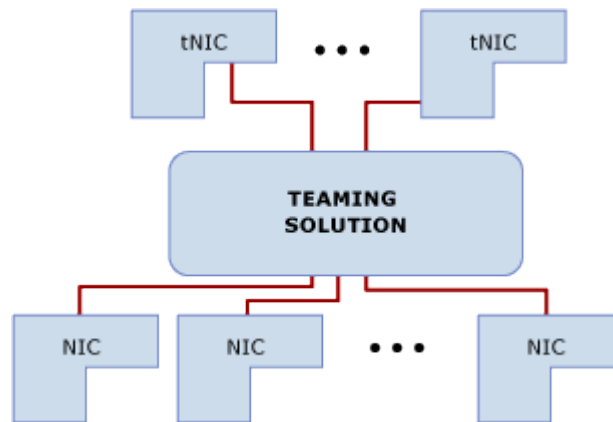


Figure 49 – Teaming solution architecture

Two or more physical network adapters are connected to the NIC Teaming solution multiplexing unit, which then presents one or more virtual adapters (also known as *team network adapters*) to the operating system. There are several different algorithms that distribute inbound and outbound traffic between the physical network adapters. In current non-Microsoft NIC Teaming solutions, the team network adapters divide traffic by virtual LAN (VLAN) so that applications can connect to different VLANs at the same time. Technically, this separation of traffic is not part of NIC Teaming. However, because other commercial implementations of NIC Teaming have this capability, the Windows Server 2012 R2 implementation also includes it.

## NIC Teaming Configurations

There are two basic sets of algorithms that are used for NIC Teaming:

- Algorithms that require the switch to participate in the teaming, also known as switch-dependent modes. These algorithms usually require all the network adapters of the team to be connected to the same switch.
- Algorithms that do not require the switch to participate in the teaming, also referred to as switch-independent modes. Because the switch does not know that the network adapter is part of a team, the team network adapters can be connected to different switches. Switch-independent modes do not require that the team members connect to different switches, they merely make it possible.

There are two common choices for switch-dependent modes of NIC Teaming:

- **Generic or static teaming (IEEE 802.3ad draft v1).** This mode requires configuration on the switch and the computer to identify which links form the team. Because this is a statically configured solution, no additional protocol assists the switch and the computer to identify incorrectly plugged cables or other errors that could cause the team to fail. This mode is typically supported by server-class switches.
- **Dynamic teaming (IEEE 802.1ax, LACP).** IEEE 802.1ax uses the Link Aggregation Control Protocol (LACP) to dynamically identify links between the computer and a specific switch. This enables the

automatic creation of a team and, in theory, the expansion and reduction of a team simply by the transmission or receipt of LACP from the peer network adapter. Typical server-class switches support IEEE 802.1ax, but most switches require manual administration to enable LACP on the port.

Both modes should result in inbound and outbound traffic approaching the practical limits of the aggregated bandwidth because the pool of links in the team functions as a single pipeline.

## Traffic Distribution Algorithms

Outbound traffic can be distributed among the available links in many ways. Try to keep all packets that are associated with a single flow (TCP stream) on a single network adapter. This is needed to minimize out-of-order packet arrival scenarios.

NIC Teaming in Windows Server 2012 R2 supports the following traffic distribution methods:

- **Hyper-V switch port.** In cases where virtual machines have independent media access control (MAC) addresses, the MAC address of the virtual machine can provide the basis for dividing traffic. There is an advantage in using this scheme in virtualization. Because the adjacent switch can determine that specific source MAC addresses are on only one connected network adapter, the switch will balance the egress load (the traffic from the switch to the computer) on multiple links, based on the destination MAC address for the virtual machine. This is particularly helpful when used with virtual machine queue. However, this mode might not be specific enough to get a well-balanced distribution, and it limits a single virtual machine to the bandwidth that is available on a single network adapter.

**Note** - Windows Server 2012 R2 uses the Hyper-V switch port as the identifier rather than the source MAC address, because in some instances, a virtual machine might be using more than one MAC address on a switch port.

- **Hashing.** This algorithm creates a hash based on components of the packet, and then it assigns packets that have that hash value to one of the available network adapters. This keeps all packets from the same TCP stream on the same network adapter. Hashing alone usually creates balance across the available network adapters. Some NIC Teaming solutions that are available on the market monitor the distribution of the traffic, and they reassign specific hash values to different network adapters in an attempt to better balance the traffic. The dynamic redistribution is known as smart load balancing or adaptive load balancing.

The components that can be used as inputs to the hashing function include the following:

- Source and destination MAC addresses
- Source and destination IP addresses, with or without considering the MAC addresses (2-tuple hash)
- Source and destination TCP ports, usually used along with the IP addresses (4-tuple hash)

The quadruple hash creates a finer distribution of traffic streams, which results in smaller streams that can be independently moved between network adapters. However, it cannot be used for traffic that is not TCP or UDP traffic or where the TCP and UDP ports are hidden from the stack, such as traffic that is protected by Internet Protocol security (IPsec). In these cases, the hash falls back to a 2-tuple hash. If the traffic is not IP traffic, the hash generator will use the source and destination MAC addresses.

- **Dynamic NIC Teaming.** Windows Server 2012 R2 uses the concept of flowlets to achieve dynamic LBFO. Flowlets are groups of TCP/IP packets that exist in most network traffic flows, and result from the inherent "burstiness" of TCP traffic. The advantage of using Flowlets to load balance is that they are smaller in size than flows and arrive more often, so enable better accuracy and quicker rebalancing of network traffic between team members.

In Windows Server 2012, flowlets are not detected and therefore rebalancing is not performed. Flowlets always follow the same path as the previous flowlet from that flow. With Windows Server 2012 R2, each flowlet is independently routed to the least used NIC in the team. Each packet within the flowlet then uses the same team member. With MAC address rewrite, the adjacent switches are unaware that flows are moving around.

## NIC Teaming In Virtual Machines

NIC Teaming in Windows Server 2012 R2 also works within a virtual machine. This allows a virtual machine to have virtual network adapters that are connected to more than one Hyper-V switch and still have connectivity even if the network adapter under that switch gets disconnected. This is particularly important when working with features such as Single Root I/O Virtualization (SR-IOV) because SR-IOV traffic does not go through the Hyper-V switch. Thus, it cannot be protected by a team that is under a Hyper-V switch. With the virtual machine teaming option, an administrator can set up two Hyper-V switches, each connected to its own SR-IOV-capable network adapter. At that point:

- Each virtual machine can then install a virtual function from one or both SR-IOV network adapters. Then, in the event of a network adapter disconnect, the virtual machine can fail over from the primary virtual function to the backup virtual function.
- Alternatively, the virtual machine might have a virtual function from one network adapter and a non-virtual function network adapter to the other switch. If the network adapter that is associated with the virtual function gets disconnected, the traffic can fail over to the other switch without loss of connectivity.

### Requirements

To fully utilize NIC Teaming, including the new Dynamic LBFO, you will require the following:

- Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2
- At least 1 network adaptor to create the team.

### Why This Matters

**NIC teaming is incredibly important for customers to ensure that network connectivity is resilient. From both the host, and the guest perspectives, the network traffic that is passing through the virtualized environment is relied upon by the rest of the business, and having a supported NIC teaming solution that is provided, in the box, manageable through both GUI and PowerShell, and provides a number of different operating modes to suit different environments, is of significant benefit to customers.**

## Failover Clustering

A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide **Cluster Shared Volume (CSV)** functionality that

provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service.

You can manage failover clusters by using the [Failover Cluster Manager](#) snap-in and the [Failover Clustering Windows PowerShell cmdlets](#). You can also use the tools in [File and Storage Services to manage file shares on file server clusters](#).

There are a number of practical applications with Failover Clustering:

- Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines.
- Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V.

With Windows Server 2012 R2, Microsoft supports the construction of Failover Clusters with up to 64 physical nodes, and from a virtualization perspective, 8,000 concurrently running virtual machines on top of the cluster. It's important to note though, that Hyper-V in Windows Server 2012 R2 supports 1,024 VMs per host, so you don't need the full complement of 64 nodes to achieve the 8,000 VMs per cluster.

From a virtualization perspective, the Failover Cluster provides the VM with high availability. If a physical host were to fail, the virtual machines running on that host would also go down. This would be a disruptive shutdown, and would incur VM downtime. However, as that physical node was part of a cluster, the remaining cluster nodes would coordinate the restoration of those downed VMs, starting them up again, quickly, on other available nodes in the cluster. This is automatic, and without IT admin intervention. This ensures that workloads running on a cluster, have a higher level of availability than those running on standalone physical servers.

In Windows Server 2008 R2 and earlier, running virtual machines on a cluster required that the VMs be placed on shared storage. Shared storage in that scenario meant a SAN, either iSCSI or FC. With Windows Server 2012 and subsequently Windows Server 2012 R2, Failover Clustering now supports the VMs being placed on a file share, accessible using the SMB 3.0 protocol, over the network. This provides administrators with considerably more flexibility when deploying their infrastructure, and also allows for a simpler deployment and management experience.

For customers who still do utilize SAN storage as their chosen shared storage solution, it is highly recommended to present LUNs from the SAN, to each node in the cluster, and enable those LUNs as **Cluster Shared Volumes**.

## Cluster Shared Volumes

Cluster Shared Volumes (CSVs) in a Windows Server 2012 R2 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume. With CSVs, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership, or dismounting and remounting a volume. CSVs also help simplify managing a potentially large number of LUNs in a failover cluster. To each of the cluster nodes in the cluster, the CSV appears as a consistent file namespace i.e. C:\ClusterStorage\Volume1.

CSVs provide a general-purpose, clustered file system which is layered above NTFS. They are not restricted to specific clustered workloads. (In Windows Server 2008 R2, CSVs only supported the Hyper-V workload) CSV applications include:

- Clustered virtual hard disk (VHD) files for clustered Hyper-V virtual machines
- Scale-out file shares to store application data for the Scale-Out File Server role. Examples of the application data for this role include Hyper-V virtual machine files and Microsoft SQL Server data

With the release of Windows Server 2012 R2, there have been a number of improvements in CSV.

### **Optimized CSV Placement Policies**

CSV ownership is now automatically distributed and rebalanced across the failover cluster nodes.

In a failover cluster, one node is considered the owner or "coordinator node" for a CSV. The coordinator node owns the physical disk resource that is associated with a logical unit (LUN). All I/O operations that are specific to the file system are through the coordinator node. Distributed CSV ownership increases disk performance because it helps to load balance the disk I/O.

Because CSV ownership is now balanced across the cluster nodes, one node will not own a disproportionate number of CSVs. Therefore, if a node fails, the transition of CSV ownership to another node is potentially more efficient.

This functionality is useful for a Scale-Out File Server that uses storage spaces because it ensures that storage spaces ownership is distributed.

In Windows Server 2012, there is no automatic rebalancing of coordinator node assignment. For example, all LUNs could be owned by the same node. In Windows Server 2012 R2, CSV ownership is evenly distributed across the failover cluster nodes based on the number of CSVs that each node owns.

Additionally in Windows Server 2012 R2, ownership is automatically rebalanced when there are conditions such as a CSV failover, a node rejoins the cluster, you add a new node to the cluster, you restart a cluster node, or you start the failover cluster after it has been shut down.

### **Increased CSV Resiliency**

Windows Server 2012 R2 includes the following improvements to increase CSV resiliency:

- Multiple Server service instances per failover cluster node. There is the default instance that handles incoming traffic from Server Message Block (SMB) clients that access regular file shares, and a second CSV instance that handles only inter-node CSV traffic. This inter-node traffic consists of metadata access and redirected I/O traffic.
- CSV health monitoring of the Server service

A CSV uses SMB as a transport for I/O forwarding between the nodes in the cluster, and for the orchestration of metadata updates. If the Server service becomes unhealthy, this can impact I/O performance and the ability to access storage. Because a cluster node now has multiple Server service instances, this provides greater resiliency for a CSV if there is an issue with the default instance. Additionally, this change improves the scalability of inter-node SMB traffic between CSV nodes.

If the Server service becomes unhealthy, it can impact the ability of the CSV coordinator node to accept I/O requests from other nodes and to perform the orchestration of metadata updates. In Windows Server 2012 R2, if the Server service becomes unhealthy on a node, CSV ownership automatically transitions to another node to ensure greater resiliency.

In Windows Server 2012, there was only one instance of the Server service per node. Also, there was no monitoring of the Server service.

### **CSV Cache Allocation**

Windows Server 2012 introduced a new feature known as CSV Cache. The CSV cache provides caching at the block level of read-only unbuffered I/O operations by allocating system memory (RAM) as a write-through cache. (Unbuffered I/O operations are not cached by the cache manager in Windows Server 2012.) This can improve performance for applications such as Hyper-V, which conducts unbuffered I/O operations



when accessing a VHD. The CSV cache can boost the performance of read requests without caching write requests. By default, the CSV cache was disabled.

In Windows Server 2012 R2, you can allocate a higher percentage of the total physical memory to the CSV cache. In Windows Server 2012, you could allocate only 20% of the total physical RAM to the CSV cache. You can now allocate up to 80%.

Increasing the CSV cache limit is especially useful for Scale-Out File Server scenarios. Because Scale-Out File Servers are not typically memory constrained, you can accomplish large performance gains by using the extra memory for the CSV cache. Also, **in Windows Server 2012 R2, CSV Cache is enabled by default.**

## Active Directory-Detached Clusters

In Windows Server 2012 R2, you can deploy a failover cluster without dependencies in Active Directory Domain Services (AD DS) for network names. This is referred to as an Active Directory-detached cluster. When you deploy a cluster by using this method, the cluster network name (also known as the administrative access point) and network names for any clustered roles with client access points are registered in Domain Name System (DNS). However, no computer objects are created for the cluster in AD DS. This includes both the computer object for the cluster itself (also known as the cluster name object or CNO), and computer objects for any clustered roles that would typically have client access points in AD DS (also known as virtual computer objects or VCOs).

**Note** - The cluster nodes must still be joined to an Active Directory domain

With this deployment method, you can create a failover cluster without the previously required permissions to create computer objects in AD DS or the need to request that an Active Directory administrator pre-stages the computer objects in AD DS. Also, you do not have to manage and maintain the cluster computer objects for the cluster. For example, you can avoid the possible issue where an Active Directory administrator accidentally deletes the cluster computer object, which impacts the availability of cluster workloads.

The option to create an Active Directory-detached cluster was not available in Windows Server 2012. In Windows Server 2012, you can only deploy a failover cluster where the network names for the cluster are in both DNS and AD DS.

An Active Directory-detached cluster uses Kerberos authentication for intra-cluster communication. However, when authentication against the cluster network name is required, the cluster uses NTLM authentication.

**Windows Server 2012 and Windows Server 2012 R2 clusters also have the ability to start up with no AD DS dependencies, which provides more flexibility for datacenters with virtualized domain controllers running on the cluster.**

## Cluster Quorum & Dynamic Witness

The quorum for a cluster is determined by the number of voting elements that must be part of active cluster membership for that cluster to start properly or continue running. Traditionally, every node in the cluster has a single quorum vote. In addition, a quorum witness (when configured) has an additional single quorum vote. With Windows Server 2012, you could configure one quorum witness for each cluster. A quorum witness could be a designated disk resource or a file share resource. Each element can cast one "vote" to determine whether the cluster can run. Whether a cluster has quorum to function properly is determined by the majority of the voting elements in the active cluster membership.

To increase the high availability of the cluster, and the roles that are hosted on that cluster, it is important to set the cluster quorum configuration appropriately.

The cluster quorum configuration has a direct effect on the high availability of the cluster, for the following reasons:

- It helps ensure that the failover cluster can start properly or continue running when the active cluster membership changes. Membership changes can occur because of planned or unplanned node shutdown, or when there are disruptions in connectivity between the nodes or with cluster storage.
- When a subset of nodes cannot communicate with another subset of nodes (a split cluster), the cluster quorum configuration helps ensure that only one of the subsets continues running as a cluster. The subsets that do not have enough quorum votes will stop running as a cluster. The subset that has the majority of quorum votes can continue to host clustered roles. This helps avoid partitioning the cluster, so that the same application is not hosted in more than one partition.
- Configuring a witness vote helps the cluster sustain one extra node down in certain configurations..

Be aware that the full function of a cluster depends on quorum in addition to the following factors:

- Network connectivity between cluster nodes
- The capacity of each node to host the clustered roles that get placed on that node
- The priority settings that are configured for the clustered roles

For example, a cluster that has five nodes can have quorum after two nodes fail. However, each remaining node would serve clients only if it had enough capacity to support the clustered roles that failed over to it and if the role settings prioritized the most important workloads.

### **Witness Configuration**

As a general rule when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster contains an even number of voting nodes, you should configure a disk witness or a file share witness. The cluster will be able to sustain one additional node down. In addition, adding a witness vote enables the cluster to continue running if half the cluster nodes simultaneously go down or are disconnected.

A disk witness is usually recommended if all nodes can see the disk. A file share witness is recommended when you need to consider multisite disaster recovery with replicated storage. Configuring a disk witness with replicated storage is possible only if the storage vendor supports read-write access from all sites to the replicated storage.

### **Node Vote Assignment**

In Windows Server 2012, as an advanced quorum configuration option, you could choose to assign or remove quorum votes on a per-node basis. By default, all nodes are assigned votes. Regardless of vote assignment, all nodes continue to function in the cluster, receive cluster database updates, and can host applications.

You might want to remove votes from nodes in certain disaster recovery configurations. For example, in a multisite cluster, you could remove votes from the nodes in a backup site so that those nodes do not affect quorum calculations. This configuration is recommended only for manual failover across sites.

The configured vote of a node can be verified by looking up the **NodeWeight** common property of the cluster node by using the **Get-ClusterNode** Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote configured. A value of 1 indicates that the quorum vote of the node is assigned, and it is managed by the cluster.

The vote assignment for all cluster nodes can be verified by using the Validate Cluster Quorum validation test.

Additional considerations:

- Node vote assignment is not recommended to enforce an odd number of voting nodes. Instead, you should configure a disk witness or file share witness. For more information, see Witness configuration in this topic.
- If dynamic quorum management is enabled, only the nodes that are configured to have node votes assigned can have their votes assigned or removed dynamically.

### Dynamic Quorum Management

In Windows Server 2012, as an advanced quorum configuration option, you can choose to enable dynamic quorum management by cluster. When this option is enabled, the cluster dynamically manages the vote assignment to nodes, based on the state of each node. Votes are automatically removed from nodes that leave active cluster membership, and a vote is automatically assigned when a node rejoins the cluster. **By default, dynamic quorum management is enabled.**

**Note** - With dynamic quorum management, the cluster quorum majority is determined by the set of nodes that are active members of the cluster at any time. This is an important distinction from the cluster quorum in Windows Server 2008 R2, where the quorum majority is fixed, based on the initial cluster configuration.

With dynamic quorum management, it is also possible for a cluster to run on the last surviving cluster node. By dynamically adjusting the quorum majority requirement, the cluster can sustain sequential node shutdowns to a single node.

The cluster-assigned dynamic vote of a node can be verified with the **DynamicWeight** common property of the cluster node by using the **Get-ClusterNode** Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote. A value of 1 indicates that the node has a quorum vote.

The vote assignment for all cluster nodes can be verified by using the Validate Cluster Quorum validation test.

Additional considerations:

- Dynamic quorum management does not allow the cluster to sustain a simultaneous failure of a majority of voting members. To continue running, the cluster must always have a quorum majority at the time of a node shutdown or failure.
- If you have explicitly removed the vote of a node, the cluster cannot dynamically add or remove that vote

### Dynamic Witness

In Windows Server 2012 R2, if the cluster is configured to use dynamic quorum (the default), the witness vote is also dynamically adjusted based on the number of voting nodes in current cluster membership. If there are an odd number of votes, the quorum witness does not have a vote. If there is an even number of votes, the quorum witness has a vote.

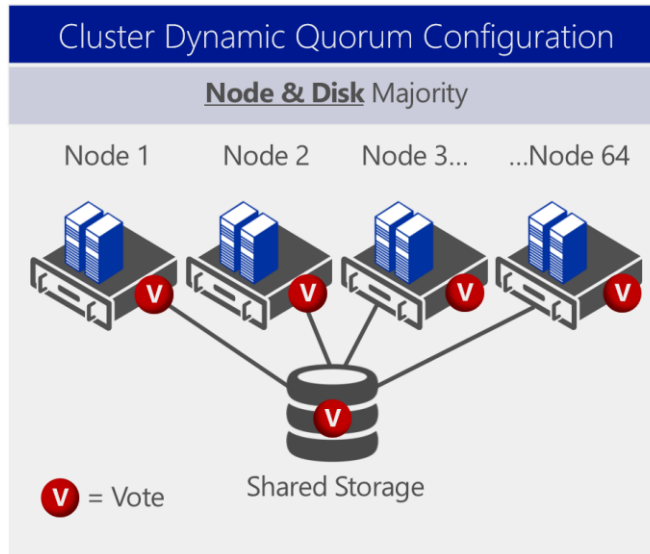


Figure 50 – 64 Node Cluster – All Nodes and Witness Disk have a Vote

As you can see from the figure above, with all being well, for a 64 node cluster, we're using the 'Node and Disk Majority' quorum configuration, that would have been automatically chosen for us as the recommended default, but that would be 64 votes – an even number. We want an odd number, hence the use of the witness disk as the 65<sup>th</sup> vote. However, if we were to lose a node, taking us down to 63 running nodes:

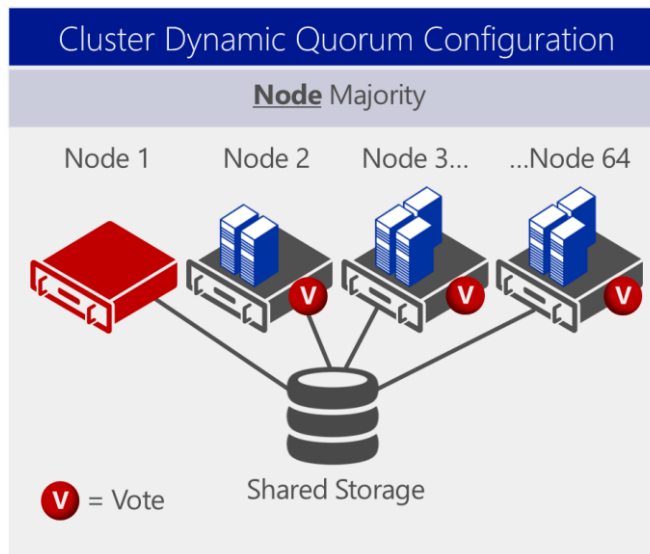


Figure 51 – 64 Node Cluster – 63 Nodes have a Vote

In this instance, our workloads have failed over to alternative nodes, and the vote from the downed node has been removed. This would leave us with 63 active nodes, each with a vote, and the witness disk, which also, previously, had a vote. This would be a total of 64 votes – an even number. We want to ensure we have an odd number of votes, as discussed earlier, and in Windows Server 2012 R2, we automatically adjust the vote of the witness disk to 0. The cluster quorum is also automatically adjusted, this time to 'Node Majority'.

The quorum witness vote is also dynamically adjusted based on the state of the witness resource. If the witness resource is offline or failed, the cluster sets the witness vote to "0."

Dynamic witness significantly reduces the risk that the cluster will go down because of witness failure. The cluster decides whether to use the witness vote based on the number of voting nodes that are available in the cluster.

This change also greatly simplifies quorum witness configuration. You no longer have to determine whether to configure a quorum witness because the recommendation in Windows Server 2012 R2 is to always configure a quorum witness. The cluster automatically determines when to use it.

**Note** - In Windows Server 2012 R2, we recommend that you always configure a quorum witness.

Windows Server 2012 R2 Preview also includes the new **WitnessDynamicWeight** cluster common property that you can use to view the quorum witness vote.

## VM Drain on Shutdown

In Windows Server 2012, if you shut down a cluster node without first draining the node, the virtual machines are put into a saved state, and then moved to other nodes and resumed. This means that there is an interruption to the availability of the virtual machines. If it takes too long to save state the virtual machines, they may be turned off, and then restarted on another node. In Windows Server 2012 R2 however, the cluster automatically live migrates all running virtual machines before shutdown.

This change provides a safety mechanism to help ensure that a server shutdown (or any action that shuts down the Cluster service) does not cause unplanned downtime for running virtual machines. This increases the availability of applications that run within the guest operating system.

We still recommend that you put a node into maintenance mode or move all virtual machines to other nodes before you shut down a cluster node. This is the safest way to drain any running clustered roles.

To enable or disable this functionality, configure the **DrainOnShutdown** cluster common property. By default, this property is enabled (set to a value of "1").

## VM Network Health Detection

In Windows Server 2012, if there is a network disconnection at the virtual machine level, the virtual machine continues to run on that computer even though the virtual machine may not be available to users.

In Windows Server 2012 R2, there is now a **protected network** check box in the virtual machine settings. If a network disconnection occurs on a protected virtual network, the cluster live migrates the affected virtual machines to a host where that external virtual network is available. For this to occur there must be multiple network paths between cluster nodes.



Figure 52 – A Protected Network for a Virtual Machine


This setting is available in the advanced features of the network adapter. By default, the setting is enabled. You can configure this setting on a per network basis for each virtual machine. Therefore, if there is a lower priority network such as one used for test or for backup, you can choose not to live migrate the virtual machine if those networks experience a network disconnection.

**Note** - If there are no available networks that connect to other nodes of the cluster, the cluster removes the node from cluster membership, transfers ownership of the virtual machine files, and then restarts the virtual machines on another node.

This change increases the availability of virtual machines when there is a network issue. If live migration occurs, there is no downtime because live migration maintains the session state of the virtual machine.

## Enhanced Cluster Dashboard

In Windows Server 2012, you had to click each failover cluster name to view status information. In Windows Server 2012 R2, Failover Cluster Manager now includes a cluster dashboard that enables you to quickly view the health status of all managed failover clusters. You can view the name of the failover cluster together with an icon that indicates whether the cluster is running, the number and status of clustered roles, the node status, and the event status.






Name	Role Status	Node Status	Event Status
 CLUS1.contoso.loc	<a href="#">0 total</a>	<a href="#">2 total</a>	 <a href="#">Critical: 3, Error: 6, Warning: 1</a>
 CLUS2.contoso.loc	<a href="#">0 total</a>	<a href="#">2 total</a>	None in the last hour

Figure 53 – An Example Cluster Dashboard

If you manage multiple failover clusters, this dashboard provides a convenient way for you to quickly check the health of the failover clusters.

## VM Monitoring

In clusters running Windows Server 2012 and Windows Server 2012 R2, administrators can monitor services on clustered virtual machines that are also running Windows Server 2012 or Windows Server 2012 R2. You can monitor any Windows service (such as SQL or IIS) in your virtual machine or any ETW event occurring in your virtual machine. When the condition you are monitoring gets triggered, the Cluster Service logs an event in the error channel on the host and takes recovery actions. These actions could be that the service is restarted, or the clustered virtual machine can be restarted or moved to another node (depending on service restart settings and cluster failover settings).

Note - You will only see services listed that run on their own process e.g. SQL, Exchange. The IIS and Print Spooler services are exempt from this rule. You can however setup monitoring for any NT service using Windows PowerShell using the **Add-ClusterVMMonitoredItem** cmdlet – with no restrictions:

```
Add-ClusterVMMonitoredItem -VirtualMachine TestVM -Service spooler
```

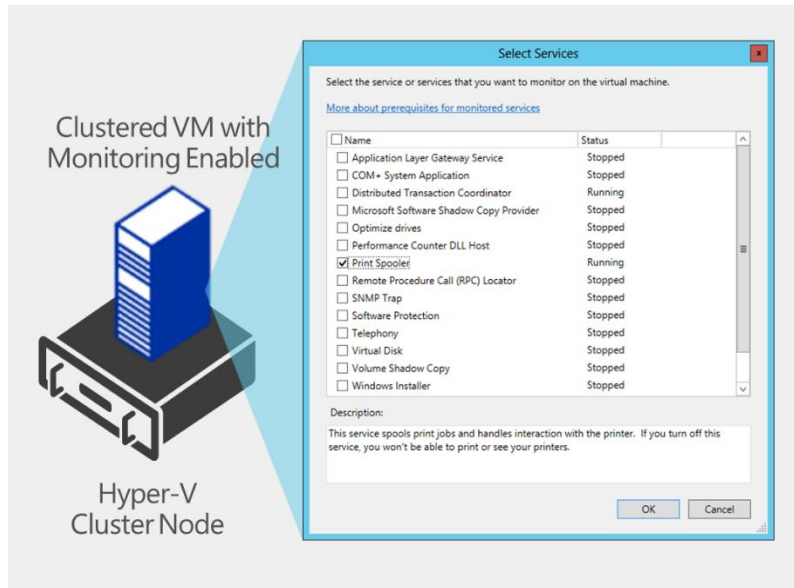


Figure 54 – VM Monitoring at the Service/App Level

When a monitored service encounters an unexpected failure, the sequence of recovery actions is determined by the Recovery actions on failure for the service. These recovery actions can be viewed and configured using Service Control Manager inside the guest. In the example below, on the first and second service failures, the service control manager will restart the service. On the third failure, the service control manager will take no action and defer recovery actions to the cluster service running in the host.

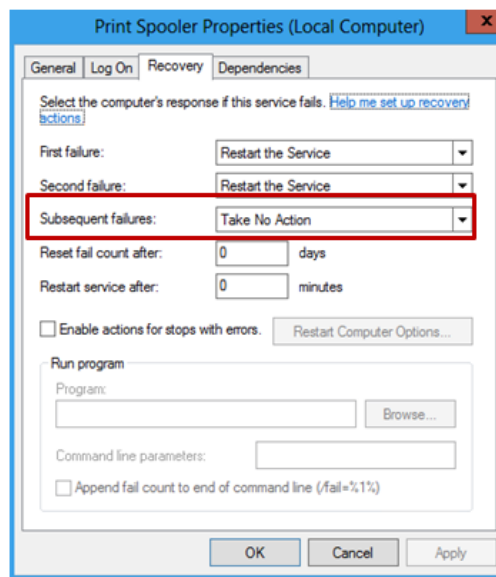


Figure 55 – Recovery Actions inside the Guest OS for Print Spooler

The cluster service monitors the status of clustered virtual machines through periodic health checks. When the cluster services determines that a virtual machine is in a “critical” state i.e. an application or service inside the virtual machine is in an unhealthy state, the cluster service takes the following recovery actions:

- Event ID 1250 is logged on the host – this could be monitored by a centralized monitoring solution such as System Center Operations Manager

- The virtual machine status in Failover Cluster Manager will indicate that the virtual machine is in an “Application Critical” state
- Recovery action is taken on the virtual machine in “Application Critical” state. Firstly, the VM is restarted on the same node. **Note** - The restart of the virtual machine is forced but graceful. On the second failure, the virtual machine restarted and failed over to another node in the cluster. **Note** - The decision on whether to failover or restart on the same node is configurable and determined by the failover properties for the virtual machine.

## Failover Priority, Affinity and Anti-Affinity

### Priority

With Windows Server 2012, and subsequently, Windows Server 2012 R2, Failover Clustering provides a new capability that allows the administrator to define the startup order of virtual machines running on that cluster, so that in the event of a failure, with a number of virtual machines needing to restart as quickly as possible, some will be prioritized over others, depending on the settings chosen.

This capability allows administrators to ensure that when resources are constrained upon failover, the most important VMs will start first and get the resources they need, before starting other, less important VMs.

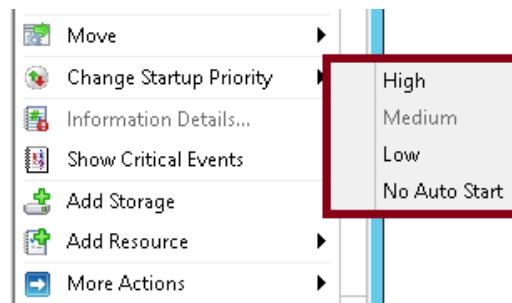


Figure 56 – Setting VM Priority for a Clustered VM

In addition, the Cluster service takes offline lower priority virtual machines when high-priority virtual machines do not have the necessary memory and other resources to start after a node failure. The freed-up resources can be assigned to high-priority virtual machines. When necessary, preemption starts with the lowest priority virtual machines and continues to higher priority virtual machines. Virtual machines that are preempted are later restarted in priority order.

### Affinity

With Windows Server 2012 and Windows Server 2012 R2 Failover Clusters, the administrator can define **preferred** and **possible owners**.

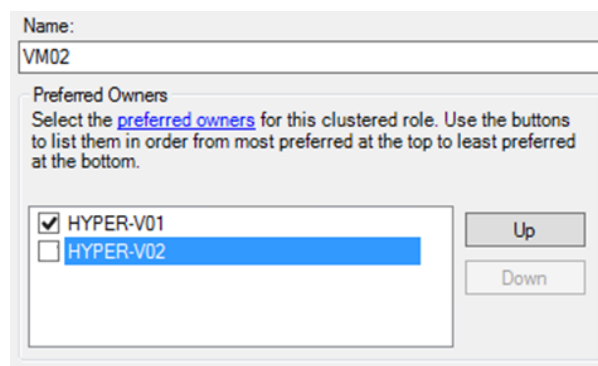


Figure 57 – Setting Preferred Ownership for a Clustered VM



For a given VM (technically any cluster Group) you can configure the preference for node order on failover. So let's say that this VM normally runs on Node A and you always want it next to go to Node C if it is available, then preferred owners is a way for you to define a preference of first go to this node, then next go to this other node, then go to this next node. It's a priority list, and clustering will walk that list in where to place the VM. This will give you more explicit control of where VMs go. More information about Preferred and Possible Owners: <http://support.microsoft.com/kb/299631>.

Possible owners on the other hand, is where, for a given VM (technically any cluster Resource) you can configure the nodes which the VM has the possibility of failing over to. By default it's all nodes, but if you have a specific node you never want this VM to failover to you can remove it from being a possible owner and prevent it.

### Anti-Affinity

Possible owners is one way to try to ensure that related VMs stay apart on different nodes, with each VM having a different set of possible owners to another, however there is another way.

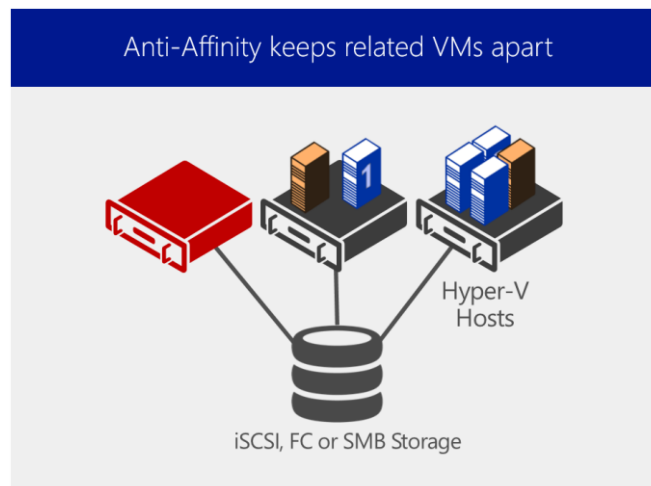


Figure 58 – Anti-Affinity can keep clustered VMs, like those in orange, apart on the cluster.

AntiAffinityClassNames is a cluster group property in Windows Failover Clustering used to identify cluster groups that should not be hosted on the same node. When working with clustered Hyper-V environments there is a 1:1 relationship between a cluster group and a virtual machine, and thus we can configure Anti-Affinity for virtual machines using the AntiAffinityClassNames property.

Once configured, Failover Clustering will attempt, as much as it can, to keep VMs that are part of the same group, on different nodes in the cluster. When combined with the failover priority and the preferred and possible ownership, the granular configuration options provided enable precise control and placement for key virtualized workloads.

### Cluster-Aware Updating

In previous releases of Windows Server, the server updating tools (e.g. WSUS) did not factor in the fact that a group of servers could be members of a highly-available cluster. As failover clusters are all about high availability of services hosted on the cluster, one would almost never patch all cluster nodes at the same time. So patching a failover cluster usually meant a fair number of manual steps, scripting/tools, and juggling administrator's attention across a number of clusters to successfully update them all during a short monthly maintenance window.

Cluster-Aware Updating (CAU) is a key capability, built into Windows Server 2012 R2 that addresses precisely this gap. CAU allows you to update clustered servers with little or no loss in availability during the update process. During an Updating Run, CAU transparently puts each node of the cluster into node maintenance mode, temporarily fails over the “clustered roles” off to other nodes, installs the updates and any dependent updates on the first node, performs a restart if necessary, brings the node back out of maintenance mode, fails back the original clustered roles back onto the node, and then proceeds to update the next node. CAU is cluster workload-agnostic, and it works great with Hyper-V, and a number of File Server workloads.

From a Hyper-V perspective specifically, CAU will work in conjunction with the Failover Cluster to Live Migrate any running virtual machines to a different physical node, to ensure no downtime for key applications and workloads running inside the VMs, whilst the host fabric is kept patched and up to date.

When the administrator triggers a CAU scan, CAU works with the nodes themselves, triggering them to perform an update check against their own update source, which could be Microsoft Update, Windows Update or a WSUS for instance.

CAU facilitates the adoption of consistent IT processes across the enterprise. Updating Run Profiles can be created for different classes of failover clusters and then managed centrally on a file share to ensure that CAU deployments throughout the IT organization apply updates consistently, even if the clusters are managed by different lines-of-business or administrators.

Updating Runs can be scheduled on regular daily, weekly, or monthly intervals to help coordinate cluster updates with other IT management processes and CAU provides an extensible architecture to update the cluster software inventory in a cluster-aware fashion. This can be used by publishers to coordinate the installation of software updates that are not published to Windows Update or Microsoft Update or that are not available from Microsoft, for example, updates for non-Microsoft device drivers.

CAU self-updating mode enables a “cluster in a box” appliance (a set of clustered physical machines running Windows Server 2012 and Windows Server 2012 R2, typically packaged in one chassis) to update itself. Typically, such appliances are deployed in branch offices with minimal local IT support to manage the clusters. Self-updating mode offers great value in these deployment scenarios.

There are two modes for CAU: self-updating mode as shown in the first figure below, and remote-updating mode as shown in the second.

In **self-updating mode** the CAU role runs on one of the nodes, which acts as the update coordinator. The update coordinator is the node that ensures that the update process is applied throughout the cluster. The CAU role is also cluster-aware, which means that more than one node can become the CAU update coordinator - but only one node at a time can be an update coordinator




Figure 59 – The CAU Update Coordinator runs on the cluster in Self-Updating Mode.

On the other hand, the **remote updating mode** allows for a remote machine running Windows Server 2012, Windows Server 2012 R2, Windows 8 or Windows 8.1 to act as the CAU update coordinator. Utilizing this method is good for those that want to see the real-time progress of the updates, and for using CAU with Windows Server 2012 / R2 Server Core OS's that require the patching.

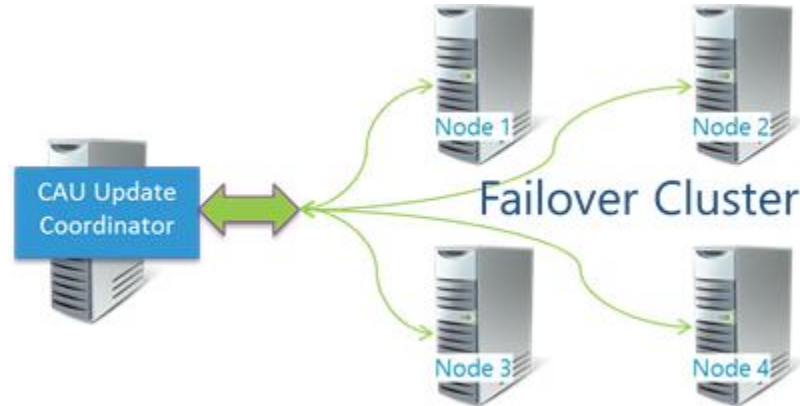


Figure 60 – The CAU Update Coordinator runs on the cluster in Self-Updating Mode.

CAU can be triggered using the GUI or you can utilize the included PowerShell cmdlets.

### Requirements

To deploy Failover Clustering, and take advantage of the other Failover Clustering capabilities, you will require the following:

- Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2
- For Cluster Shared Volumes, you will require shared storage, either iSCSI or Fibre Channel
- For VM Monitoring, you will need to make the appropriate firewall exceptions and correct domain configuration
- For Cluster Aware Updating, you will require a WSUS infrastructure, or connection from the cluster nodes, to the internet to access Microsoft/Windows Update

### Why This Matters

**Cluster-Aware Updating provides a significant benefit to organizations looking to streamline and automate the maintenance of the Windows Server clusters, and is of particular importance with larger scale Hyper-V clusters. CAU facilitates the adoption of consistent IT processes across the enterprise. Updating Run Profiles can be created for different classes of failover clusters and then managed centrally on a file share to ensure that CAU deployments throughout the IT organization apply updates consistently, even if the clusters are managed by different lines-of-business or administrators. Ensuring the hosts are patched and up to date, in accordance with internal policy is important, but ensuring this is performed without causing downtime to the virtualized workloads on top, is perhaps of even greater importance. This built-in capability ensures customers of all shapes and sizes can automate and orchestrate the centralized patching of their Windows Server 2012 and 2012 R2 clusters.**

## Guest Clustering

With Windows Server 2012 Hyper-V, Microsoft provided full support for the virtualization of VMs that themselves, would form part of a cluster, at the guest OS level. An example would be a clustered SQL AlwaysOn configuration, which itself, would require multiple nodes, all being virtual machines, and would also require access to some shared storage. The configuration could look very similar to the following figure:




Figure 61 – A Guest Cluster running on a Hyper-V Cluster

In the above example, we have a simple 3 node Hyper-V physical cluster, and on top, we have a 2 node Guest Cluster, built from 2 virtual machines, both of which have direct access to some form of shared storage. Guest clusters that require shared storage, on the Microsoft platform, can take advantage of multiple storage choices. For instance, if a customer is using SMB storage as their main Hyper-V cluster repository, that same SMB storage can be provisioned, over the network, and exposed to the VMs themselves. In the same way, VMs, through their virtual network adaptors, can be given access to iSCSI storage. Both of these scenarios however, would require exposing the VMs, via their virtual network adaptors, to the underlying storage fabric directly, instead of using VHDs or VHDX files for the shared storage.

With Windows Server 2012, Virtual Fibre Channel was also introduced. As discussed earlier, this presents FC storage directly through to the VM, or VMs, again, allowing the construction of guest clusters with access to shared storage.

These guest cluster configuration, as stated earlier, are fully supported by Microsoft, and in addition, can be combined with features such as Live Migration, and Dynamic Memory, meaning customers can virtualize their clustered workloads, without sacrificing the key features for density and agility. Also, guest clusters can benefit significantly from the failover priority, affinity and anti-affinity features discussed earlier, to ensure that guest cluster nodes stay positioned optimally in respect of each other, and the underlying physical hosts.

The advantage of a guest cluster, is the second level of resiliency. Should a physical host fail, only a subset of the guest cluster nodes would fail too, and the application-level resiliency provided by the guest cluster would pick up the workload quickly.




Figure 62 – A Guest Cluster running on a Hyper-V Cluster that has a single node failure

In this example, the physical node has failed, and the VM that was previously running on there, also went down, but the physical Hyper-V cluster ensures that VM restarts quickly, and on a different node to the one remaining guest cluster VM. The application-level resiliency ensures that on the whole, the application or workload was only down for a very short period of time.

The challenge however, is that in these configurations, the underlying storage (FC, iSCSI, SMB) was exposed to the user of a virtual machine. In private or public cloud deployments, there is often the need to hide the details of the underlying fabric from the user or tenant administrator.

### Shared VHDX

In Windows Server 2012 R2, you can now share a virtual hard disk file (in the .vhdx file format) between multiple virtual machines. You can use these .vhdx files as shared storage for a virtual machine failover cluster, or guest cluster. For example, you can create shared .vhdx files for data disks and for the disk witness. (You would not use a shared .vhdx file for the operating system virtual hard disk.)

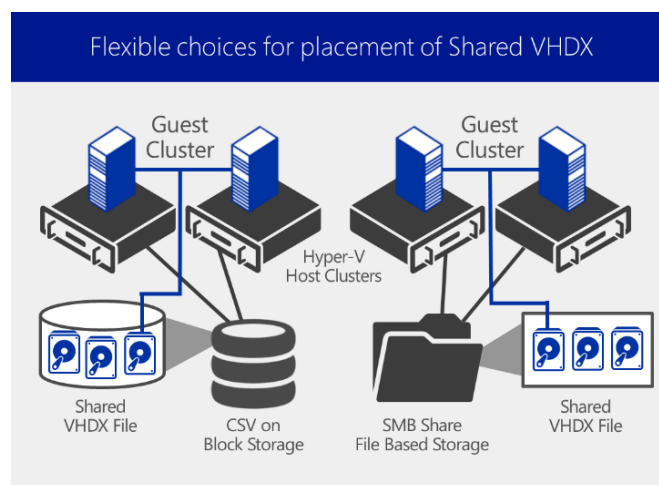


Figure 63 – A Guest Cluster using Shared VHDX on CSV/SMB Storage

This change also enables easier deployment of guest cluster configurations. A shared .vhdx file configuration is easier to deploy than solutions like virtual Fibre Channel or iSCSI. When you configure a virtual machine to use a shared .vhdx file, you do not have to make storage configuration changes such as zoning and LUN masking, and none of the underlying storage infrastructure is exposed to the users of the virtual machine.

Using a shared virtual hard disk is ideal for the following situations:

- SQL Server database files.
- File server services running within a virtual machine.
- Database files that reside on shared disks

Shared virtual hard disk functionality in guest failover clusters exclusively uses the .vhdx file format. Although the shared virtual hard disk must use the .vhdx file format for the data drive, the operating system disk for a virtual machine can use either the .vhd or the .vhdx file format.

A Hyper-V guest failover cluster that uses a shared virtual hard disk has two preferred deployment models. The shared virtual hard disk for the guest failover cluster can be deployed on:

- Cluster Shared Volumes (CSVs) on block storage (including clustered Storage Spaces).
- Scale-Out File Server with SMB 3.0 on file-based storage.

A Hyper-V guest failover cluster can be configured and deployed by using Hyper-V Manager and Failover Cluster Manager. You can also use Windows PowerShell.

## Requirements

To deploy a Guest Cluster using Shared VHDX, you will require the following:

- 2 x Windows Server 2012 R2 with Hyper-V, or Hyper-V Server 2012 R2 cluster nodes
- Servers must belong to the same Active Directory domain.
- Availability of configured shared storage resources—for example, CSVs on block storage (such as clustered storage spaces) or a Scale-Out File Server cluster (running Windows Server 2012 R2) with SMB 3.0 (for file-based storage).
- Sufficient memory, disk, and processor capacity within the failover cluster to support multiple virtual machines that are implemented as guest failover clusters

## Why This Matters

**Guest Clustering is becoming increasingly important as customers look to achieve even higher levels of availability for their key workloads. Workloads that have a clustering capability as part of their core, such as SQL server, in certain configurations, rely on shared storage presented through to the OS running that workload. In a virtualized environment, this typically is very restrictive and not always feasible. With Windows Server 2012 R2, support for guest clusters is flexible in terms of the presentation of storage, and of the Hyper-V features that are still accessible for those VMs once part of a Guest Cluster – something that competitive platforms can't offer.**

**Customers have the flexibility to connect their shared storage, via iSCSI, Virtual Fibre Channel, or SMB 3.0, straight to their VMs, ensuring they can provide the extra level of application resiliency they demand, whilst maximizing the investment in existing storage, but for**

customers who don't wish to expose the underlying storage fabric to the virtual guests themselves, the Shared VHDX capability abstracts the underlying storage and presents a high-performance shared virtual hard disk, which still resides on that shared storage, through to multiple guest for flexibility, and secure abstraction.

## Incremental Backup

In Windows Server 2008 R2 and earlier, backing up data required you to perform full file backups. This meant that you had to either back up the virtual machine and snapshots as flat files when offline, or use Windows Server or third party tools to back up the virtual machine itself with a normal backup of the operating system and data. **Windows Server 2012 R2 supports incremental backup of virtual hard disks while the virtual machine is running.**

Incremental backup of virtual hard disks lets you perform backup operations more quickly and easily, saving network bandwidth and disk space. Because backups are VSS aware, hosting providers can run backups of the Hyper-V environment, backing up tenant virtual machines efficiently and offering additional layers of service to customers without the need for a backup agent inside the virtual machines.

Incremental backup can be independently enabled on each virtual machine through the backup software. Windows Server 2012 R2 uses "recovery snapshots" to track the differences between backups. These are similar to regular virtual machine snapshots, but they are managed directly by Hyper-V software. During each incremental backup, only the differences are backed up (note the blue highlights in the following figure).

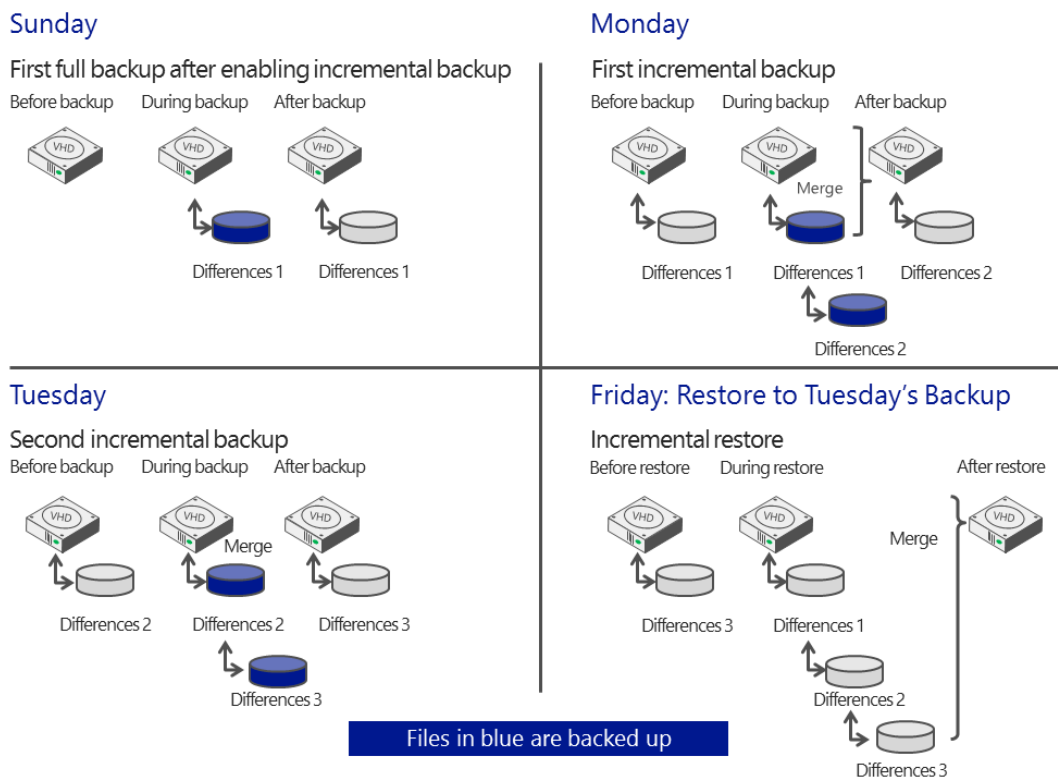


Figure 64 – Example of incremental backup of a virtual hard disk

The figure above illustrates the incremental backup of a virtual machine with one virtual hard disk and shows 3 days of backups (Sunday, Monday, and Tuesday) followed by a restore (Friday). Of note in this example are the following points:

- To enable change tracking, the virtual machine must be configured to use incremental backup, and a full backup must be performed after incremental backup is enabled (see Sunday).
- During an incremental backup, the virtual machine will briefly be running off of two levels of recovery snapshots. The earlier recovery snapshot is merged into the base virtual hard disk at the end of the backup process.
- The virtual machine's configuration XML files are very small and are backed up often. For simplicity, they are not shown in the figure above

## Windows Azure Backup Integration

Windows Azure Backup is a cloud-based backup solution that enables server data to be backed up to and recovered from an off-premises datacenter (the cloud) to help protect against data loss and corruption. To reduce storage and bandwidth utilization, Windows Azure Backup performs block-level incremental backup. To increase security, the data is compressed and encrypted before leaving the server.

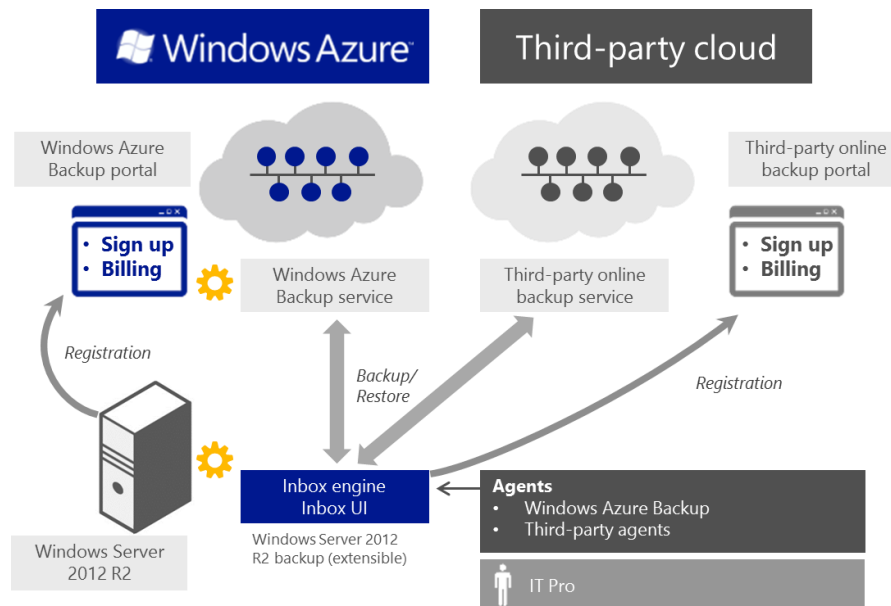


Figure 65 – Windows Azure Backup

Once a customer has completed the appropriate registration, the Windows Azure Backup agent is installed on the target server, which provides both the MMC interface and PowerShell interface for management. Once configuration is completed, the administrator can quickly and easily register the target server, and define the backup schedule, along with which files and folders to protect.

Windows Azure Backup is a very valuable addition to a local backup policy. It would allow a customer to recover data in case of disasters (server destroyed/stolen, disk crash), or alternatively, recover data in case of data loss scenarios such as data accidentally deleted, volume deleted, viruses. It also serves as a valid low-cost alternative to long-term tape archiving.

There are a number of target segments that Windows Azure Backup is aimed at:

- Small Business: Low-cost backup & recovery solution for single server backups.



- Departmental Backups: Low-cost backup alternative for departments in mid to large sized organizations
- Remote office backup and recovery consolidation: Consolidate backups of remote offices.

### Requirements

To deploy a Windows Server Backup, with integration with Windows Azure Backup, you will require the following:

- Windows Server 2012 R2 (with Hyper-V if protecting VMs) or Hyper-V Server 2012 R2
- Windows Azure Backup subscription/billing account
- Windows Azure Backup Agent installed and configured, along with appropriate certificates

### Why This Matters

**Having an easy to use, powerful and intelligent backup capability built into Windows Server 2012 R2, enables customers to efficiently protect their important workloads, so much so, that in the event of a failure, customers can easily restore from these backups, and be up and running again quickly. Add to this, the ability to integrate the Windows Server Backup, with Windows Azure Backup, and customers can move away from tape for long term archiving, or off-site backup, and instead backup to Azure, taking advantage of Azure's limitless storage capacity. Customers can flexibly pay as they consume. The Windows Azure Backup integration is perfect for smaller organizations, departments, or remote sites, and, in the event of a complete site loss, would allow streamlined and simple recovery to get that data back, and the workloads up and running as quickly as possible.**

## Hyper-V Replica

Business continuity depends on fast recovery of business functions after a downtime event, with minimal or no data loss. There are number of reasons why businesses experience outages, including power failure, IT hardware failure, network outage, human errors, IT software failures, and natural disasters. Depending on the type of outage, customers need a high availability solution that simply restores the service.

However, some outages that impact the entire datacenter, such as a natural disaster or an extended power outage, require a disaster recovery solution that restores data at a remote site and brings up the services and connectivity. Organizations need an affordable and reliable business continuity solution that helps them recover from a failure.

Beginning with Windows Server 2008 R2, Hyper-V and Failover Clustering could be used together to make a virtual machine highly available and minimize disruptions. Administrators could seamlessly migrate virtual machines to a different host in the cluster in the event of outage or to load balance their virtual machines without impacting virtualized applications.

While these measures could protect virtualized workloads from a local host failure or scheduled maintenance of a host in a cluster, they did not protect businesses from outages of an entire datacenter. While Failover Clustering can be used with hardware-based SAN replication across datacenters, these are typically expensive. Hyper-V Replica, a key feature of Windows Server 2012 R2, now offers an affordable in-box disaster recovery solution.

Hyper-V Replica provides asynchronous replication of virtual machines for the purposes of business continuity and disaster recovery. This asynchronous replication, in Windows Server 2012 R2, is now configurable. The administrator has the choice of:

- 30 seconds
- 5 minutes
- 15 minutes

Hyper-V Replica is incredibly simple to enable, through a wizard in Hyper-V Manager, through PowerShell, or through System Center Virtual Machine Manager. Once Replica is enabled for a particular virtual machine, the initial replication can begin.

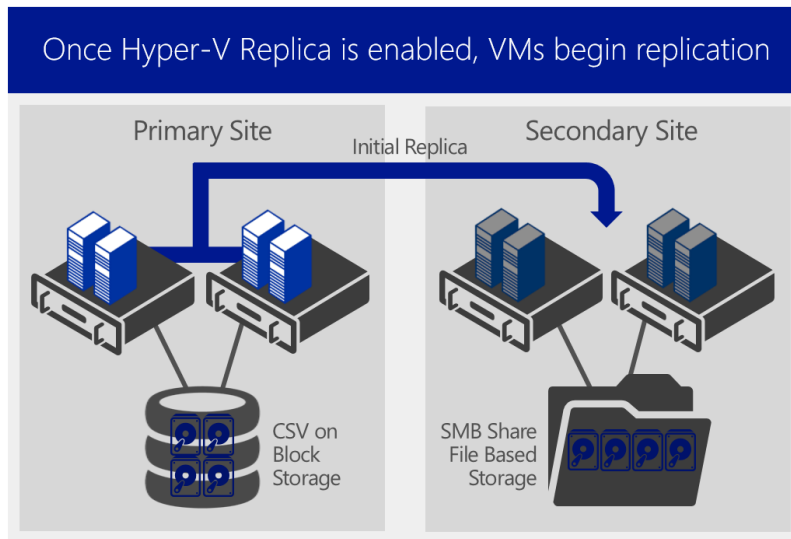


Figure 66 – Hyper-V Replica: Initial Replication

The initial replication can be triggered immediately, scheduled for a later time, or even exported to a USB drive, for physical transportation to the target site, before replication begins. If a customer already has a backup of the source VM on the target site, this can also be used as the replication target.

As you can see from the figure above, Replica provides complete flexibility for replication. Being software based, there is no requirement on specific hardware on either site, ensuring complete flexibility and low cost. Administrators also have the ability to specify additional recovery points, outside of just the most recent. These recovery points, in Windows Server 2012 R2, are configurable up to a 24 hour period. The administrator also has the flexibility to choose what is replicated. For instance, if a VM had 4 virtual disks, but only 3 had important data, the 4<sup>th</sup> could be excluded from the replication, saving valuable bandwidth and disk space.

Hyper-V Replica tracks the write operations on the primary virtual machine and replicates these changes to the Replica server efficiently over a WAN.

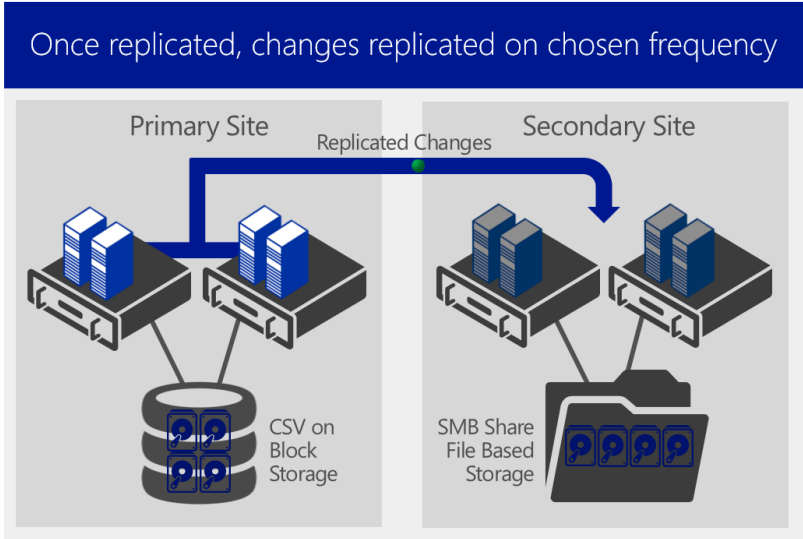


Figure 67 – Hyper-V Replica: Replicated Changes

The network connection between the two servers uses the HTTP or HTTPS protocol and supports both Windows-integrated and certificate-based authentication. For an encrypted connection, you should choose certificate-based authentication. Hyper-V Replica can also be closely integrated with Windows Failover Clustering, and provides easier replication across different migration scenarios in the primary and Replica servers. As it is integrated with Failover Clustering, Hyper-V Replica has full understanding of Live Migration, ensuring that VMs that are moving around the clustered environments, will still be replicated to their target sites as appropriate.




Figure 68 – Hyper-V Replica: VMs started on second site after failure

In the event of a disaster, VMs can be quickly and easily started on the second site, ensuring minimal data loss, and downtime for key applications and workloads.

### Extended Replication

In Windows Server 2012, Hyper-V Replica would allow replication every 5 minutes, and only between 2 points. So, for instance, a customer could replicate their VMs to a Service Provider, but that would be the

further that the VM could be replicated. The Service Provider wouldn't easily be able to replicate your VM on to a DR site of their own, for instance.

With Windows Server 2012 R2 Hyper-V, not only have the replication intervals become configurable by the administrator, with the choice of 30 seconds, 5 minutes, or 15 minutes, but the replication capabilities have been enhanced to allow for replication of a VM to a tertiary location.

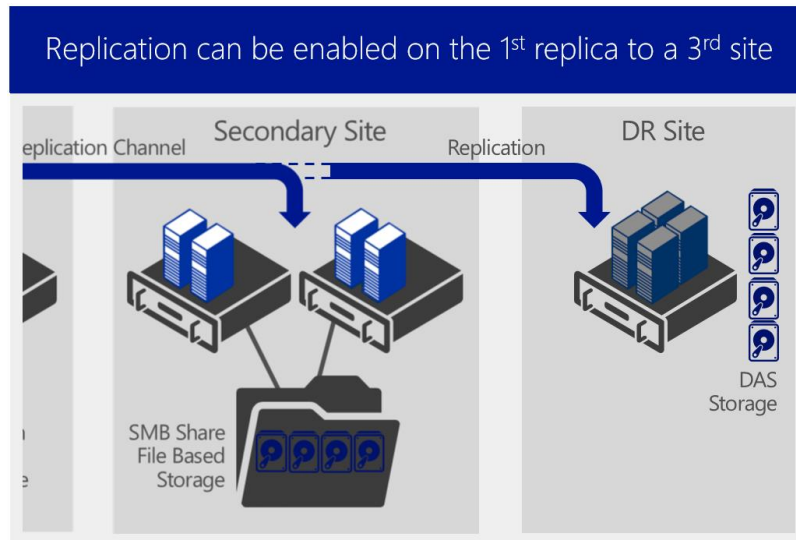


Figure 69 – Hyper-V Replica: Extended Replication

In the figure above, the VMs that were previously replicated to a second site, have now been replicated again, to a third site, providing an extra level of resiliency and peace of mind for the customer. The replication intervals for the extended replication are either 5 minutes, or 15 minutes. The data that was replicated from Primary to Secondary, will be the same data that will be replicated from Secondary to Tertiary, however the administrator has granular control over ports, initial replication (from Secondary to Tertiary), and recovery points. Again, as stated earlier, this provides complete flexibility, agnostic of hardware, with the above figure replicating to a low-cost DR site using DAS storage.

Again, this can be configured through Hyper-V Manager, PowerShell or SCVMM.

### Requirements

To take advantage of Hyper-V Replica, and the Extended Replication, you will need the following:

- At least 3 x Windows Server 2012 R2 with Hyper-V or Hyper-V Server 2012 R2

### Why This Matters

**Hyper-V Replica is a powerful, in-box engine for replication. Its ease of use, simple configuration, and hardware-agnostic approach to VM replication make it applicable to businesses large and small. The new improvements in Replica, such as the increased flexibility for replication time, down to near-synchronous levels at 30 second intervals, and up to 15 minutes, give customers more flexibility depending on their specific requirements. In addition, being able to be managed through a simple GUI, or through PowerShell, means**

customers have the flexibility to setup, configure and manage Replica in the way that's right for them.

With Extended Replication, customers can replicate not only between their own sites, but also to a 3<sup>rd</sup> site, ensuring that if there was a catastrophic failure of their 2 datacenters, the 3<sup>rd</sup> place would have the data they need, and would be able to restore those VM's quickly and efficiently. This built-in replication capability removes cost-barriers for adoption and enables customers of all shapes and sizes to benefit quickly.

## Windows Azure Hyper-V Recovery Manager

Hyper-V Replica is a replication engine, to replicate VMs from site to site (and again, to site), and is configured, typically, using Hyper-V Manager, PowerShell, or VMM. In smaller environments, customers can quickly configure Replica, on a VM by VM basis, and can maintain simple and efficient control of this replicating environment whilst it stays small. As the environment grows however, and increases in complexity, it becomes more difficult to configure and manage Hyper-V Replica without PowerShell, or automation tools such as System Center Orchestrator, both of which could be used to automate the failover of replicated VMs from site to site.

Windows Azure Hyper-V Recovery Manager however, provides another way.




Figure 70 – Windows Azure Hyper-V Recovery Manager

Windows Azure Hyper-V Recovery Manager can help you protect important services by coordinating the replication and recovery of System Center managed private clouds at a secondary location.

System Center Virtual Machine Manager Private Clouds can be protected through automating the replication of the virtual machines that compose them at a secondary location. The ongoing asynchronous replication of each VM is provided by Hyper-V Replica and is monitored and coordinated by Windows Azure Hyper-V Recovery Manager. The replication **is not to** Windows Azure. The replication exists only between

the source and destination sites. Windows Azure Hyper-V Recovery Manager only has visibility into the metadata of the VMs, in order to fail them over successfully.

The service helps automate the orderly recovery in the event of a site outage at the primary datacenter. VMs can be brought up in an orchestrated fashion to help restore service quickly. It can also incorporate specific manual activities, and scripts into the process, to ensure granular control over the failover process. This process can also be used for testing recovery, or temporarily transferring services.

## Requirements

To take advantage of Windows Azure Hyper-V Recovery Manager, you will need the following:

- A Windows Azure account that has the Windows Azure Recovery Services feature enabled.
- Enable the windows azure recovery services preview, see enable windows azure preview features.
- At least two System Center Virtual Machine Manager servers running System Center 2012 SP1 (For Windows Server 2012 Hyper-V hosts) or System Center 2012 R2, located in different datacenters.
- At least one cloud configured on the source VMM server you want to protect, and one cloud on the destination VMM server that you want to use for protection and recovery.
- One or more virtual machines located in the source cloud that you want to protect.
- Verify that VM networks are configured on the source VMM server you want to protect, and that corresponding VM networks are created on the destination VMM server. Ensure that the networks are connected to appropriate source and destination clouds.
- A management certificate (.cer and .pfx files) that you will upload to the Hyper-V Recovery Vault

## Why This Matters

**Hyper-V Replica, on its own, is a replication engine, which with the inbox MMC interfaces, enables easy configuration for a small number of VMs. If however, you have a significant number of VMs to configure and replicate, this could be time consuming without scripting. In addition, upon failover, how can you orchestrate specific VMs failing over before others, or in groups, without scripting?**

**Windows Azure Hyper-V Recovery Manager wrappers Hyper-V Replica with a powerful orchestration capability that is defined and controlled from the Cloud, specifically, Windows Azure. This means that customers can quickly and easily define recovery plans, connect their System Center Virtual Machine Manager-managed-sites to Windows Azure Hyper-V Recovery Manager, and use the web-based Azure portal to control failover between sites, with no investment in additional infrastructure needed on each site, or additional sites creating.**

**This enables customers to realize value quickly, without needing to deploy complex appliances on both sites, lowering the cost to implement and decreasing time to value.**

# Virtualization Innovation

We've spent considerable time discussing the 4 key areas around scalability, performance and density, security and multi-tenancy, flexibility, and finally, high availability and resiliency. As you can see, Windows Server 2012 R2 is delivering innovation in all of those areas, ensuring Hyper-V is the best platform for your key workloads going forward, offering the most compelling features and capabilities, at the lowest cost.

There are however, a number of capabilities within Windows Server 2012 R2 that push beyond just virtualization. Features and capabilities that aim to change the way customers virtualize. To push the boundaries in terms of performance, usability, and flexibility. To close out the whitepaper, we'll focus on 3 of those key capabilities.

## Generation 2 VMs

Virtual machine generation determines the virtual hardware and functionality that is presented to the virtual machine. In Windows Server 2012 R2 Hyper-V there are two supported virtual machine generations, generation 1 and generation 2. Generation 2 virtual machines have a simplified virtual hardware model, and it supports Unified Extensible Firmware Interface (UEFI) firmware instead of BIOS-based firmware. Additionally, the majority of legacy devices are removed from generation 2 virtual machines.

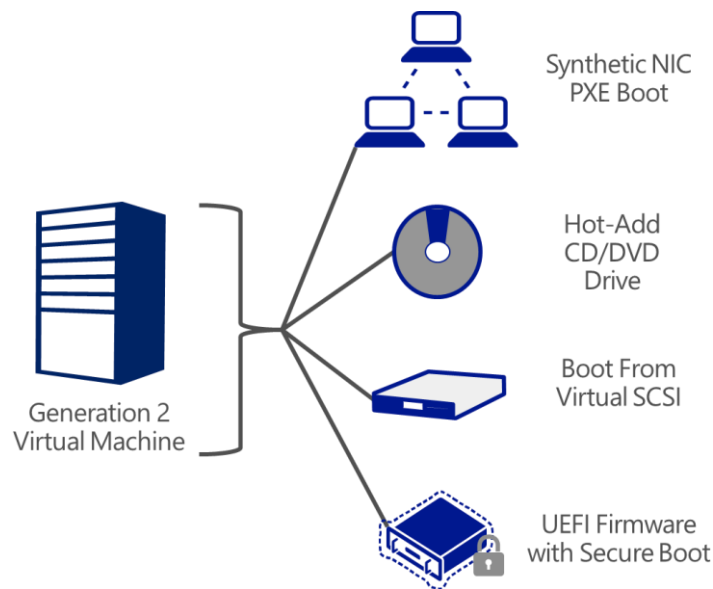


Figure 71 – Key Changes with Generation 2 VMs

**Generation 2** - Provides the following new functionality on a virtual machine:

- **PXE boot by using a standard network adapter** - In previous versions of Hyper-V, if you wanted to perform a remote installation of the guest operating system by using PXE boot, you were required to install a legacy network adapter for the PXE boot, in addition to the standard network adapter that you would use after the operating system was installed.

Generation 2 virtual machines support PXE boot by using a standard network adapter, so there is no need to install a legacy network adapter. The legacy network adapter has been removed from generation 2 virtual machines.

- **Boot from SCSI controller** - In previous versions of Hyper-V, you could not boot a virtual machine from a SCSI-attached virtual hard disk or from a DVD. Generation 2 virtual machines can boot from a virtual

hard disk or DVD that is attached to the SCSI controller. The virtual IDE controller has been removed from generation 2 virtual machines.

- **Secure Boot** - Secure Boot is a feature that helps prevent unauthorized firmware, operating systems, or UEFI drivers (also known as option ROMs) from running at boot time.

In addition, through adopting the generation 2 virtual machines, customers can expect to see faster boot and OS install times.

### Requirements

The following hosts and guest operating systems are supported as generation 2 virtual machines.

- Windows Server 2012 R2 Hyper-V or Hyper-V Server 2012 R2
- Windows Server 2012 or Windows Server 2012 R2
- 64-bit versions of Windows 8 or Windows 8.1

### Why This Matters

**Generation 2 VMs provide a more advanced virtual machine hardware platform for your key workloads. Whilst restricted to only supporting the most recent of Guest operating systems, the payoff is a more secure virtualized machine, with a removed reliance on emulated drivers within the guest. Customers who use Generation 2 VMs can gain additional performance flexibility, such as in PXE booting scenarios, and will see faster OS install and boot times as a result.**

## Enhanced Session Mode

To enhance the user experience when administering Hyper-V, Hyper-V and the Virtual Machine Connection tool now support redirection of local resources to a virtual machine session. This feature provides similar type of device redirection to a virtual machine as you get with a Remote Desktop Connection session.

In previous versions of Hyper-V the Virtual Machine Connection utility only provided redirection of the virtual machine screen, keyboard, and mouse along with limited copy / paste functionality. To get additional redirection abilities a Remote Desktop Connection to the virtual machine could be initiated, but this would require a network path to the virtual machine.

Starting with Hyper-V in Windows Server 2012 R2, Hyper-V can now redirect local resources to a virtual machine session through **Virtual Machine Connection tool**. The enhanced session mode connection uses a Remote Desktop Connection session via the virtual machine bus (VMBus), so no network connection to the virtual machine is required.

The following local resources can be redirected when using the Virtual Machine Connection tool.

- Display configuration
- Audio
- Printers
- Clipboard
- Smart cards
- USB devices



- Drives
- Supported Plug and Play devices

This feature is enabled by default in Client Hyper-V and is disabled by default on Hyper-V running on Windows Server




Figure 72 – Enhanced Session Mode

Enhanced session mode can be useful in the following scenarios:

- Troubleshooting a virtual machine without the need for a network connection to the virtual machine.
- Login to the virtual machine via smart card
- Printing from a virtual machine to a local printer
- Developers can now fully test and troubleshoot applications running in a virtual machine that require USB and sound redirection without the need to use Remote Desktop Connection

### Requirements

The following hosts and guest operating systems are supported with Enhanced Session Mode.

- Windows Server 2012 R2 Hyper-V or Hyper-V Server 2012 R2
- 64-bit version of Windows 8.1

### Why This Matters

Enhanced Session mode provides a much richer user experience for IT admins who are administering Hyper-V. From copy and paste, through to RDP over VMBus, these changes make Hyper-V administration more flexible, and less time consuming, whether you're running Hyper-V in the server environment, or on the desktop client within Windows 8.1.

## Automatic Virtual Machine Activation

Automatic Virtual Machine Activation (AVMA) acts as a proof-of-purchase mechanism, helping to ensure that Windows products are used in accordance with the Product Use Rights and Microsoft Software License Terms.

AVMA lets you install virtual machines on a properly activated Windows server without having to manage product keys for each individual virtual machine, even in disconnected environments. AVMA binds the virtual machine activation to the licensed virtualization server and activates the virtual machine when it starts up. AVMA also provides real-time reporting on usage and historical data on the license state of the virtual machine. Reporting and tracking data is available on the virtualization server.

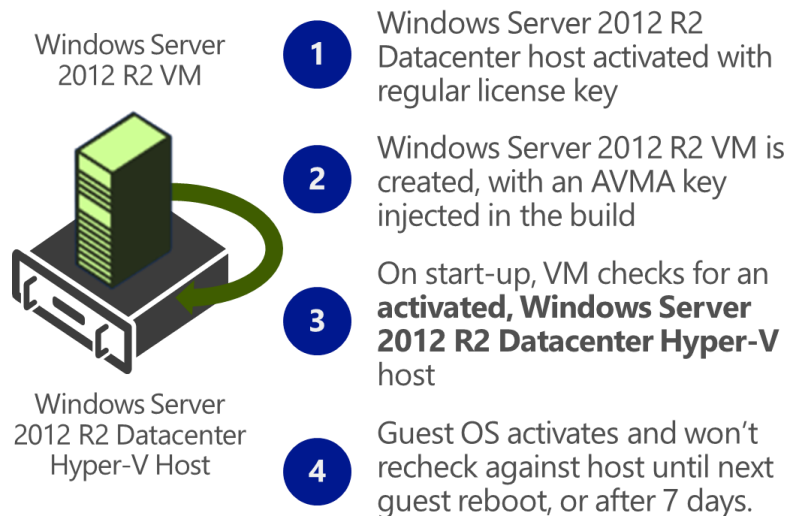


Figure 73 – Automatic VM Activation

On virtualization servers that are activated using Volume Licensing or OEM licensing, AVMA offers several benefits.

- Server datacenter managers can use AVMA to do the following:
- Activate virtual machines in remote locations
- Activate virtual machines with or without an internet connection
- Track virtual machine usage and licenses from the virtualization server, without requiring any access rights on the virtualized systems

There are no product keys to manage and no stickers on the servers to read. The virtual machine is activated and continues to work even when it is migrated across an array of virtualization servers.

Service Provider License Agreement (SPLA) partners and other hosting providers do not have to share product keys with tenants or access a tenant's virtual machine to activate it. Virtual machine activation is transparent to the tenant when AVMA is used. Hosting providers can use the server logs to verify license compliance and to track client usage history.

The registry (KVP) on the virtualization server provides real-time tracking data for the guest operating systems. Because the registry key moves with the virtual machine, you can get license information as well. By default the KVP returns information about the virtual machine, including the following:

- Fully qualified domain name
- Operating system and service packs installed

- Processor architecture
- IPv4 and IPv6 network addresses
- RDP addresses

### Requirements

To take advantage of AVMA you will require the following:

- Windows Server 2012 R2 Datacenter with Hyper-V Host.
- The guest virtual machine operating system must be Windows Server 2012 R2 Datacenter, Windows Server 2012 R2 Standard, or Windows Server 2012 R2 Essentials

### Why This Matters

**AVMA provides significant advantages to Service Providers and disconnected environments who struggle to manage license activation. Being able to securely activate the Guest OS's, against a legitimately activated host, ensures that VMs themselves don't need to be necessarily exposed to the corporate network for activation, and at the same time, customers don't need to manually activate each VM they deploy. This not only ensures legal compliance from a licensing standpoint, but also simplifies licensing management.**

# Conclusion

In this paper, we have looked at a significant number of the new capabilities that are available within Windows Server 2012 R2 Hyper-V, across 5 key investment areas:

**Scalability, Performance & Density** – We've shown how, with Hyper-V customers can run the biggest, most powerful virtual machines, to handle the demands of their biggest workloads. We saw that as hardware scale grows, customers wishing to take advantage of the largest physical systems to drive the highest levels of density, and reduce overall costs, can do so successfully with Hyper-V. In addition, we looked at a number of the points of integration between Hyper-V and hardware, driving the highest levels of performance for enterprise applications.

**Security & Multitenancy** – We looked not only at the different capabilities such as BitLocker, which enables a level of physical security for the virtualized hosts, but also some of the in-box, granular networking security capabilities now built into the Hyper-V Extensible Switch, which enables customers to securely and easily isolate and control access to their key workloads inside the virtualized environment.

**Flexible Infrastructure** – We discussed how, in a modern datacenter, customers are looking to be agile, in order to respond to changing business demands quickly, and efficiently, and how through the new innovation in Live Migration, Hyper-V provides this in the box. Being able to move workloads flexibly around the infrastructure is of incredible importance, and in addition, customers want to be able to choose where best to deploy their workloads based on the needs of that workload specifically, and to do that, Network Virtualization plays a significant part. Also, for customers with heterogeneous infrastructures, with a mixture of both Linux and Windows-based workloads, Hyper-V provides the ideal platform through the continued engineering and development to improve Linux performance on Hyper-V.

**High Availability & Resiliency** – As customers' confidence in virtualization grows, and they virtualize their more mission-critical workloads, the importance of keeping those workloads continuously available grows significantly. With Windows Server 2012 R2, having capabilities built into the platform that not only help keep those workloads highly available, but also, in the event of a disaster, quick to restore in another geographical location, is of immense importance when choosing a platform for today's modern datacenter. We discussed specific improvements at both the fabric and the workload level that help to ensure that customers can keep their most mission critical applications and data as continuously available as possible.

**Virtualization Innovation** – Finally, we discussed some of the key features that take Hyper-V beyond just virtualization. Features such as the Automatic VM Activation – a fantastic addition for Service Providers and disconnected environments where VM activation has been a challenge in days gone by. In addition, Generation 2 VMs herald a new direction for Hyper-V VMs, exposing new capabilities, for both performance, flexibility and security, now through to the virtual machines themselves.

Across each of these 5 key areas, there have been a number of improvements. Building on the innovative, solid platform of Windows Server 2012, Windows Server 2012 R2 provides customers with an enterprise-class virtualization platform, with the key features that enterprises are looking for to virtualize their mission-critical applications and workloads, but at a fraction of the cost of competing technologies.