



304 N Main St
Manheim, PA
Ph: 717-665-2141

ElectSys Reseller and Support Vendor of Windows 2012
Hyper-V in the Lancaster PA area.



Why Hyper-V?

Competitive Advantages of
Windows Server 2012 R2 Hyper-V
over VMware vSphere 5.5

October 2013 v1.0

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

- Beyond Virtualization..... 3
 - Before Windows Server 2012 R2.....3
 - Windows Server 2008 R2 Hyper-V Enhancements.....4
 - Windows Server 2012 Hyper V and Windows Server 2012 R2.....4
- Why Hyper-V?..... 5
 - Scalability, Performance & Density.....5
 - Enhanced Storage Capabilities8
 - Enhanced Networking Performance.....10
 - Enhanced Resource Management.....12
 - Virtualized Workload Performance14
 - Security & Multitenancy..... 19
 - The Hyper-V Extensible Switch.....19
 - Extending the Extensible Switch.....20
 - Physical Security.....22
 - Flexible Infrastructure.....22
 - Linux Support on Hyper-V23
 - Virtual Machine Live Cloning.....24
 - Virtual Machine Mobility24
 - Hyper-V Network Virtualization26
 - High Availability & Resiliency.....30
 - NIC Teaming.....30
 - Failover Clustering.....31
 - Guest Clustering.....32
 - Incremental Backup.....36
 - Hyper-V Replica.....37
 - Windows Azure Hyper-V Recovery Manager.....37
 - Virtualization Innovation39
 - Generation 2 VMs.....40
 - Enhanced Session Mode.....41
 - Automatic Virtual Machine Activation.....42
- Conclusion..... 44

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 competitive advantages that Windows Server 2012 R2 Hyper-V provides, over the VMware vSphere 5.5 release, focusing on key capabilities across scalability and performance, security and multitenancy, flexibility and high availability & resiliency.

Before Windows Server 2012 R2

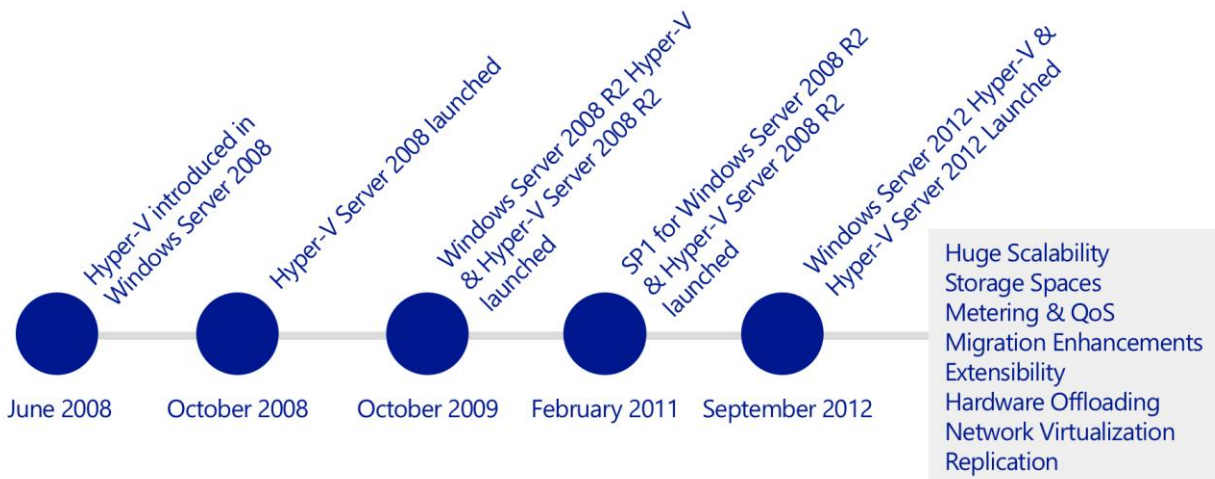


Figure 1 – Timeline of Windows Server Hyper-V & Hyper-V Server Releases

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 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, and how it compares and contrasts with vSphere 5.5. 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.

Why Hyper-V?

Virtualization technologies help customers' lower costs and deliver greater agility and economies of scale. Either as a stand-alone product or an integrated part of Windows Server, Hyper-V is a leading virtualization platform for today and the transformational opportunity with cloud computing.

With Hyper-V, it is now easier than ever for organizations to take advantage of the cost savings of virtualization, and make the optimum use of server hardware investments by consolidating multiple server roles as separate virtual machines that are running on a single physical machine. Customers can use Hyper-V to efficiently run multiple operating systems, Windows, Linux, and others, in parallel, on a single server. Windows Server 2012 R2 extends this with more features, greater scalability and further inbuilt reliability mechanisms.

In the data center, on the desktop, and now in the cloud, the Microsoft virtualization platform, which is led by Hyper-V and surrounding System Center management tools, simply makes more sense and offers better value for money when compared to the competition.

This paper will focus on comparing Windows Server 2012 R2 Hyper-V, with the standalone VMware vSphere Hypervisor, also known as ESXi, and vSphere 5.5.

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 and 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 R2 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
Host	Logical Processors	64	320	5x
	Physical Memory	1TB	4TB	4x
	Virtual CPUs per Host	512	2,048	4x
VM	Virtual CPUs per VM	4	64	16x
	Memory per VM	64GB	1TB	16x
	Active VMs per Host	384	1,024	2.7x
	Guest NUMA	No	Yes	-
Cluster	Maximum Nodes	16	64	4x
	Maximum VMs	1,000	8,000	8x

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.

Windows Server 2012 R2 Hyper-V also 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.

How does VMware Compare?

The table below shows a comparison between Windows Server 2012 R2 Hyper-V, and both the VMware vSphere Hypervisor; VMware's free standalone hypervisor, and VMware vSphere 5.5 Enterprise Plus; VMware's per-CPU licensed, most advanced edition.

	Resource	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Host	Logical Processors	320	320	320
	Physical Memory	4TB	4TB	4TB
	Virtual CPUs per Host	2,048	4,096	4,096
VM	Virtual CPUs per VM	64	8	64
	Memory per VM	1TB	1TB	1TB
	Active VMs per Host	1,024	512	512
	Guest NUMA	Yes	Yes	Yes
Cluster	Maximum Nodes	64	N/A	32
	Maximum VMs	8,000	N/A	4,000

The table above shows that Hyper-V has a number of advantages from a scalability perspective, especially when it comes to comparison with the vSphere Hypervisor. vSphere 5.5 brought a number of scalability increases for vSphere environments, doubling the number of host logical processors supported from 160 to 320, and doubling the host physical memory from 2TB to 4TB, but this still only brings vSphere up to the level that Hyper-V has been offering since September 2012, at the launch of Windows Server 2012 Hyper-V.

VMware positions the vSphere Hypervisor as simple, entry-level solution designed to allow users to experience the benefits of VMware's virtualization platform at no cost, however on closer examination, certain restrictions are imposed which prevent customers utilizing the solution at scale, meaning customers have to purchase, at significant cost, one of the more advanced vSphere editions.

Since the launch of vSphere 5.0, in 2011, VMware has regularly discussed the inclusion of 32 virtual processors within a virtual machine, yet at the time, this was exclusive to the Enterprise Plus edition of vSphere, and not the vSphere Hypervisor, vSphere 5.0 Essentials, Essentials Plus, Standard, and Enterprise editions, which were all capped at 8 virtual processors per virtual machine. With vSphere 5.1, and subsequently, 5.5, however, the Enterprise edition can now support VMs with up to 32 vCPUs, and the Enterprise Plus edition, 64 vCPUs. Compare this with Hyper-V in Windows Server 2012 and 2012 R2, and customers not only receive up to 64 virtual processors per virtual machine, but this comes with no SKU-specific restrictions. Customers are free to run the most demanding of their workloads on Hyper-V, without additional costs or expensive edition upgrades. The table also shows that both Windows Server 2012 R2 Hyper-V and vSphere 5.5 deliver up to 1TB of memory to an individual virtual machine. Previously, the vSphere Hypervisor was physically limited from consuming more than 32GB of memory, which severely restricted VM sizes, however this restriction has been lifted with the 5.5 release. From an individual host perspective, Hyper-V also supports double the number of active virtual machines per host, than both the vSphere Hypervisor and vSphere 5.5 Enterprise Plus, ensuring customers can realize even greater levels of density for their key workloads, whilst achieving a better return on investment.

Whilst virtualization itself is an incredibly important aspect within the datacenter, resiliency and high availability of workloads is of equal importance. The inclusion of Failover Clustering with Windows Server 2012 R2 enables customers to achieve massive scale with an unparalleled number of nodes within a cluster, and virtual machines per cluster. Unfortunately, the vSphere Hypervisor alone doesn't provide any high availability, or resiliency features, and customers must purchase vSphere 5.5 to unlock these features, and even then, cluster sizes are restricted to only 32 nodes, and 4,000 virtual machines per cluster, which is considerably smaller than the 64 nodes, and 8,000 VMs supported by Windows Server 2012 R2.

Enhanced Storage Capabilities

Windows Server 2012 and subsequently, 2012 R2 Hyper-V also introduce a number of enhanced storage capabilities to support the most intensive, mission-critical of workloads. These capabilities include:

- **Virtual Fiber Channel** – Enables virtual machines to integrate directly into Fiber Channel Storage Area Networks (SAN), unlocking scenarios such as fiber channel-based Hyper-V Guest Clusters.
- **Support for 4-KB Disk Sectors in Hyper-V Virtual Disks.** Support for 4,000-byte (4-KB) disk sectors lets customers take advantage of the emerging innovation in storage hardware that provides increased capacity and reliability.
- **New in R2 - Storage Spaces with Tiering** - Storage Spaces enables you to virtualize storage by grouping industry-standard disks into storage pools, and then create virtual disks called *storage spaces* from the available capacity in the storage pools. These pools now support a mix of HDD and SSD, providing a tiered pool, where hot data will reside on SSD and cold data on HDD. Fully supported as a repository for Hyper-V VMs.
- **Data Deduplication** - Windows Server 2012 R2 also provides an inbox deduplication capabilities which utilizes sub-file variable-size chunking and compression to considerably reduce storage consumption for files and folders hosted on deduplicated Windows Server volumes. With Windows Server 2012 R2, support has been added for VDI deployments. Deduplication rates for VDI deployments can range as high as 95% savings and that includes VDI deployments that utilize differencing disks for rapid provisioning.
- **New Virtual Hard Disk Format.** This new format, called VHDX, is designed to better handle current and future workloads and addresses the technological demands of an enterprise's evolving needs by increasing storage capacity, protecting data, improving quality performance on 4-KB disks, and providing additional operation-enhancing features. The maximum size of a VHDX file is 64TB.
- **Offloaded Data Transfer (ODX).** With Offloaded Data Transfer support, the Hyper-V host CPUs can concentrate on the processing needs of the application and offload storage-related tasks to the SAN, increasing performance.
- **Online Checkpoint Merge.** With the online checkpoint merge capability, customers who have taken checkpoints (snapshots), for a running virtual machine, no longer have to power down the virtual machine in order to merge the checkpoint back into the original virtual disk file, ensuring virtual machine uptime is increased and the administrator gains increased flexibility.
- **New in R2 - Online Virtual Disk Resize.** With the online virtual disk resize, administrators can grow and shrink virtual disks that are attached to a VM's virtual SCSI controller, providing an administrator with greater flexibility to respond to changing business needs.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
iSCSI/FC Support	Yes	Yes	Yes
Network File System Support	Yes (SMB 3.0)	Yes (NFS)	Yes (NFS)
Virtual Fiber Channel	Yes	Yes	Yes
3rd Party Multipathing (MPIO)	Yes	No	Yes (VAMP)
Native 4-KB Disk Support	Yes	No	No
Storage Virtualization	Yes (Spaces)	No	Yes (vSAN)
Storage Tiering	Yes	No	Yes

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Data Deduplication	Yes	No	No
Maximum Virtual Disk Size	64TB VHDX	62TB VMDK	62TB VMDK
Maximum Pass Through Disk Size	256TB+	64TB	64TB
SAN Offload Capability	Yes (ODX)	No	Yes (VAAI)
Online Checkpoint Merge	Yes	Yes	Yes
Online Virtual Disk Resize	Yes	Grow Only	Grow Only

As shown in the table, Windows Server 2012 R2 Hyper-V provides a significant number of advantages over both the vSphere Hypervisor and vSphere 5.5 Enterprise Plus. Customers building virtualized infrastructures today require the highest levels of availability and performance, and wish to maximize the investment in their chosen technologies to help drive their business forward. With Microsoft, the ability to utilize Device Specific Modules, also known as DSMs, produced by storage vendors, in conjunction with the Multipath I/O framework within Windows Server, ensures that customers run their workloads on an optimized storage configuration from the start, as the storage vendor intended, providing the highest levels of performance and availability. This framework is built into the Windows Server platform, at no cost. Unfortunately, the vSphere Hypervisor doesn't provide the ability to utilize these storage vendor specific optimizations, and in fact, only the Enterprise and Enterprise Plus editions of vSphere 5.5, through a feature known as 'vStorage APIs for Multipathing', provide this capability, meaning customers have to upgrade to higher, more costly editions in order to unlock the best performance from their storage investments.

When implementing a virtualized infrastructure, customers today look to the future to understand new technology trends and innovations that are coming down the line. One of those innovations is the rapidly emerging Advanced Format Disks, which have a 4KB physical sector size. These disks bring an increase in performance, and are natively supported by Windows Server 2012 R2 Hyper-V, but unfortunately, are not supported with the vSphere Hypervisor and vSphere 5.5, restricting future hardware upgrades.

When it comes to storage, whilst both vendors offer solutions that integrate with existing storage investments, such as those based on iSCSI or FC, both Microsoft and VMware are pushing forward with software-defined storage solutions. Microsoft, with Windows Server 2012 R2, now brings to market its 2nd release of the Storage Spaces technology that aggregates physical disks into pools of storage, and then slices them into spaces, for formatting with a file system, and placement of VMs on top, typically accessed over SMB 3.0. With 2012 R2, Tiering was introduced for Spaces, which allows the aggregation of HDD and SSD within a Storage Pool, and hot blocks are automatically moved to the SSD tier to drive increased performance.

With VMware on the other hand, they are only now testing, in beta as of October 2013, their new vSAN technology, which aggregates physical disks within each of the vSphere hosts, and allows the creation of datastores on top. One key difference with vSAN is its use of a distributed RAID cluster architecture, which does allow the creation of highly available storage across physical hosts – something that Spaces does not provide, yet at this time, vSAN is not ready for production use.

As customers introduce larger, more powerful workloads into their virtual environments, the amount of data associated with these workloads, over time, will grow. Fortunately, Windows Server 2012 R2 Hyper-V supports the creation of virtual disks, quickly and efficiently, of up to 64 Terabytes (TB) in size, allowing huge databases, file repositories or document archives to be stored within individual disks. Due to the architecture of VMFS, the VMware Virtual Machine File System, the maximum size of a VMFS datastore is still 64TB, and if it was filled with a 64TB VMDK, there would be no room remaining for certain management tasks, such as snapshots, thus 62TB is the maximum virtual disk size, an increase from 2TB in the 5.1 release. With Microsoft however, Windows Server 2012 R2 supports NTFS volumes significantly bigger than 64TB, in fact, up to 256TB, significantly larger and more

flexible than that offered by VMware. If customers do choose to implement RDMs (Raw Device Mappings) as an alternative to VMFS/VMDKs, 64TB is the maximum supported size, however with Microsoft, Windows Server 2012 R2 places no specific maximum on the size of a pass through disk. The maximum size of a physical disk attached to a virtual machine is ultimately determined by what the guest operating system supports, with more recent Windows Server operating systems supporting individual disk sizes of over 256TB. This ensures that the largest data-driven workloads can be virtualized on Hyper-V with ease.

We mentioned earlier, a capability known as 3rd Party Multipathing, and how this enables customers to optimize their Host-to-SAN integration and connectivity, maximizing their investment in both of these key elements of the virtualized infrastructure, and providing the highest levels of performance and availability for their critical workloads. Offloaded Data Transfer (ODX), a key capability of Windows Server 2012 R2 Hyper-V, is another of those features that enables organizations to maximize their investment in their current technologies. By integrating Windows Server 2012 R2 Hyper-V with an ODX-capable storage array, many of the storage-related tasks that would normally use valuable CPU and network resources on the Hyper-V hosts, are offloaded to the array itself, executing much faster, increasing performance significantly, and unlocking extra resources on the hosts themselves. VMware offer a similar capability, known as vStorage APIs for Array Integration, VAAI, but unfortunately, this capability is only available in the Enterprise and Enterprise Plus editions of vSphere 5.5, meaning customers, again, have to upgrade to higher editions to achieve higher performance from their hardware investments.

From a storage flexibility perspective also, Hyper-V within Windows Server 2012 R2 also brings additional capability to help the administrator effectively manage a changing environment. For instance, a VM is provisioned with a data virtual disk of 80GB, however over time, capacity of that disk decreases. Rather than take the workload down to add additional storage capacity, the new Online Virtual Disk Resize capability allows the IT admin to increase the size of the virtual machine disk whilst the virtual machine is running, with no downtime to the workload. Should the opposite happen, and a VM has been provisioned with disks that are too large, again, with Hyper-V within Windows Server 2012 R2, IT administrators can quickly and easily remove excess capacity from a virtual machine disk, with no downtime for the workload. With VMware vSphere 5.5 however, IT admins can grow the VMDK whilst the VM is running, but are unable to shrink the virtual machine disk.

Enhanced Networking Performance

Windows Server 2012 R2 Hyper-V also includes a number of performance enhancements within the networking stack to help customers virtualize their most intensive network workloads. These capabilities include:

- **Dynamic Virtual Machine Queue** – DVMQ dynamically distributes incoming VM 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
- **IPsec Task Offload** - IPsec Task Offload in Windows Server 2012 R2 leverages the hardware capabilities of server NICs to offload IPsec processing. This reduces the CPU overhead of IPsec encryption and decryption significantly. In Windows Server 2012 R2, IPsec Task Offload is extended to Virtual Machines as well. Customers using VMs who want to protect their network traffic with IPsec can take advantage of the IPsec hardware offload capability available in server NICs, thus freeing up CPU cycles to perform more application-level work and leaving the per packet encryption/decryption to hardware.
- **SR-IOV** - When it comes to virtual networking, a primary goal is native I/O throughput. Windows Server 2012 R2 provides the ability to assign SR-IOV functionality from physical devices directly into virtual machines. This gives VMs the ability to bypass the software-based Hyper-V Virtual Network Switch, and more directly address the NIC. As a result, CPU overhead and latency is reduced, with a corresponding rise in throughput. This is all available, without sacrificing key Hyper-V features such as virtual machine Live Migration.

- New in R2 – Virtual Receive Side Scaling** - Prior to 10GbE networking, one modern processor was usually more than enough to handle the networking workload of a VM. With the introduction of 10GbE NICs, the amount of data being sent to and received from a VM exceeded what a single processor could effectively handle. In the physical host, this challenge had a solution, namely, 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. With Windows Server 2012 R2 however, 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.




Figure 2 – vRSS enables all vCPUs to be utilized to process network traffic

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Dynamic Virtual Machine Queue	Yes	NetQueue	NetQueue
IPsec Task Offload	Yes	No	No
SR-IOV with Live Migration	Yes	No	No
Virtual Receive Side Scaling	Yes	Yes (VMXNet3)	Yes (VMXNet3)

Whilst VMware provide a capability known as NetQueue, in VMware’s own documentation, ‘Performance Best Practices for VMware vSphere 5.5’, it is noted that “On some 10 Gigabit Ethernet hardware network adapters, ESXi supports NetQueue, a technology that significantly improves performance of 10 Gigabit Ethernet network adapters in virtualized environments”. What does this mean for customers who have servers that don’t have 10GbE? With Windows Server 2012 R2 Hyper-V, and DVMQ, customers with existing 1GbE **and** 10GbE adaptors can flexibly utilize these advanced capabilities to improve performance and throughput, whilst reducing the CPU burden on their Hyper-V hosts.

When it comes to network security, specifically IPsec, VMware offers no offloading capabilities from the virtual machine through to the physical network interface, thus in a densely populated environment, valuable host CPU cycles will be lost to maintain the desired security level. With Windows Server 2012 R2 Hyper-V, the IPsec Task Offload capability will move this workload to a dedicated processor on the network adaptor, enabling customers to make dramatically better use of the resources and bandwidth that is available.

As stated earlier, when it comes to virtual networking, a primary goal is native I/O. With SR-IOV, customers have the ability to directly address the physical network interface card from within the virtual machine, reducing CPU overhead and latency whilst increasing throughput. In vSphere 5.1, VMware first introduced SR-IOV support, however, then, and now with 5.5, SR-IOV in vSphere requires the vSphere Distributed Switch – a feature only found in the highest vSphere edition, meaning customers have to upgrade to take advantage of this higher levels of performance. Also, VMware's implementation of SR-IOV unfortunately doesn't support other features such as vMotion, High Availability and Fault Tolerance, meaning customers who wish to take advantage of higher levels of performance, must sacrifice agility and resiliency. Prior to vSphere 5.1 and 5.5, VMware provided a feature that offered a similar capability to SR-IOV, and continues to offer this in 5.5. DirectPath I/O, a technology which binds a physical network card to a virtual machine, offers that same enhancement, to near native performance, however, unlike SR-IOV in Windows Server 2012 R2 Hyper-V, a virtual machine with DirectPath I/O enabled is restricted to that particular host, unless the customer is running a certain configuration of Cisco UCS. Other caveats include:

- Small Hardware Compatibility List
- No Memory Overcommit
- No vMotion (unless running certain configurations of Cisco UCS)
- No Fault Tolerance
- No Network I/O Control
- No VM Snapshots (unless running certain configurations of Cisco UCS)
- No Suspend/Resume (unless running certain configurations of Cisco UCS)
- No Endpoint Security support
- No NSX Virtualization Support

Whilst DirectPath I/O may be attractive to customers from a performance perspective, VMware ask customers to sacrifice agility, losing vMotion in most cases, and scale, having to disable memory overcommit, along with a number of other vSphere features.

No such restrictions are imposed when using SR-IOV with Windows Server 2012 R2 Hyper-V, ensuring customers can combine the highest levels of performance with the flexibility they need for an agile, scalable infrastructure.

Finally, when it comes to in-guest network performance, the inclusion of vRSS support in Windows Server 2012 R2 enhances the performance of network-intensive workloads running in VMs, by spreading the processing of traffic across multiple virtual processors and subsequently, physical processors. Outside of the VMXNET3 network device, vSphere does not provide this functionality through to virtual machines, so customers have to use that specific vNIC to see benefit.

Enhanced Resource Management

Windows Server 2012 R2 Hyper-V also includes a number of enhanced resource management capabilities that help customers to optimize the utilization of the virtualized infrastructure to drive higher levels of performance. These capabilities include:

- **Dynamic Memory Improvements** - These improvements dramatically increase virtual machine consolidation ratios and improve reliability for restart operations that can lead to lower costs, especially in environments, such as VDI, that have many idle or low-load virtual machines. Administrators can now more flexibly manage memory through the use of a Startup, Minimum and Maximum configuration option, along with the ability to adjust the memory values whilst the VM is running, increasing flexibility for the administrator. Windows Server 2012 R2 Hyper-V also includes a capability known as Smart Paging, which provides a more reliable and robust solution for VM restarts when memory is under contention.

- **Resource Metering** - 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. Resource Metering captures metrics across CPU, Memory, Disk and Network.
- **Network Quality of Service** - QoS provides the ability to programmatically adhere to a service level agreement (SLA) by specifying the minimum bandwidth that is available to a virtual machine or a port. It prevents latency issues by allocating maximum bandwidth use for a virtual machine or port.
- **New in R2 – Storage Quality of Service** – 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.




Figure 3 – Storage QoS in Action: Disabled on the left, and then enabled on the right

How does VMware Compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Dynamic Memory	Yes	Yes	Yes
Resource Metering	Yes	Yes	Yes
Network Quality of Service	Yes	No	Yes
Storage Quality of Service	Yes	No	Yes

As shown in the table, when it comes to memory management, Windows Server 2012 R2 Hyper-V, along with the VMware vSphere Hypervisor and vSphere 5.5, all provide techniques to better utilize virtual machine memory, increase density and maximize return on investment, however Microsoft’s approach to memory management is different to that of VMware. VMware claim, that through their 4 memory management techniques; Memory Ballooning, Transparent Page Sharing, Compression and Swapping, they can provide a virtual machine density greater than that of Hyper-V, yet in reality, this is false. All 4 of these memory management techniques only operate when the host is under memory pressure, heavily laden, as a reactive measure. With technologies such as Transparent Page Sharing (TPS), with the majority of hardware platforms now supporting higher performance 2MB Large Page Tables by default (LPT), TPS is unable to deduplicate memory pages as easily as it would, prior to LPT, thus the capability becomes significantly less useful. Under memory pressure, the vSphere Hypervisor host will break down large memory pages into smaller, 4KB pages, which it can then deduplicate, freeing up memory, but unfortunately, this process doesn’t occur without a cost to already limited, host performance. With compression and swapping, whilst both help to keep virtual

machines operable, it's too little too late, with performance of key workloads at this point, becoming severely degraded.

With Dynamic Memory, Hyper-V works intuitively with the guest operating system, delivering, and reclaiming memory from the virtual machine in a way that is optimal for the guest operating system, ensuring resources are provided appropriately, and a consistent level of performance is achieved for key workloads, ultimately providing the highest levels of density, and the greatest return on investment. With Windows Server 2012 R2 Hyper-V and Dynamic Memory, administrators also have increased flexibility to make runtime changes to the virtual machine memory to meet changing demands of the workloads.

As we move towards more cloud-oriented infrastructures, especially in multi-tenanted environments, hosting providers and enterprises must be able to measure the amount of data center resources (compute, network, and storage) that are consumed by each workload. These can be used to charge external customers (known as chargeback), or for internal accounting (known as showback) for cross-departmental budget management scenarios of an enterprise. Resource Metering, which is a standard feature of Windows Server 2012 R2 Hyper-V, when combined with new performance counters, exposes a wealth of information from which chargeback and showback models can be built. While the VMware vSphere Hypervisor, and vSphere 5.5 both enable the capturing of information within vCenter, organizations must purchase, at additional cost to vSphere 5.5, vCenter Chargeback Manager, which is only available as part of the vCenter Operations Management Suite Enterprise, in order to utilize the information in a meaningful manner.

Whilst chargeback and showback are two important elements for a private cloud, ensuring service levels are met is equally important, whether the primary business is that of a hosting provider, serving external customers, or an enterprise organization, serving internal business units with chargeable resources. Either way, ensuring the highest levels of performance is imperative, and with Windows Server 2012 Hyper-V, Quality of Service (QoS), both for networking and storage control, is a standard feature, enabling organizations to ensure that Service Level Agreements (SLAs) for key workloads are met, and at the same time, intensive virtual machines don't consume more than their allocated allowance. With VMware however, QoS, or I/O Control, is only available in the Enterprise Plus edition of vSphere 5.5, so for those customers who wish to implement stringent SLAs, customers must upgrade, at additional cost, to VMware's highest edition.

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.

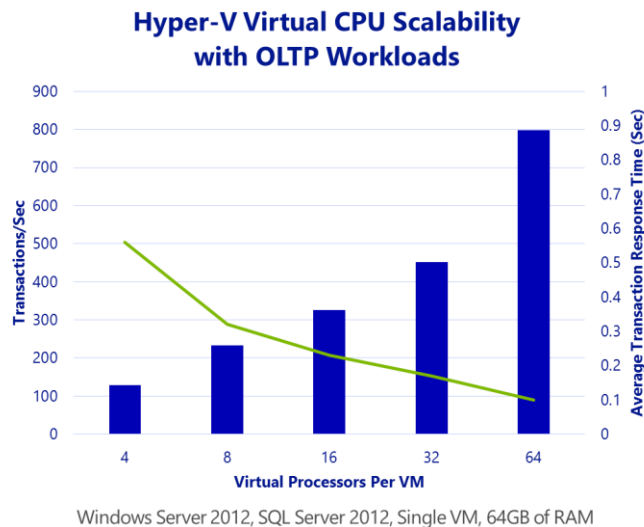


Figure 4 – Graph of a Hyper-V Virtual CPU Scalability with OLTP Workloads

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:




Figure 5 – 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 6 – 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:

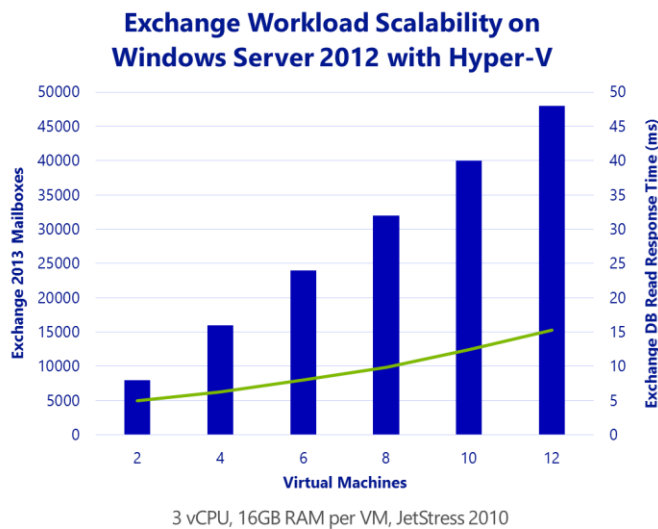


Figure 7 – Graph of virtualized Exchange 2013 scalability

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:




Figure 8 – 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 9 – 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.

SAP on Hyper-V

One of the most common workloads within enterprise environments is SAP Enterprise Resource Planning (ERP); a solution that provides access to critical data, applications, and analytical tools, and helps organizations

streamline processes across procurement, manufacturing, service, sales, finance, and HR. With a demanding workload like SAP ERP, many customers' natural thought process may be to run the solution on physical servers, and there are a significant number of existing SAP benchmarks that highlight huge scale and performance on a physical platform.

On June 24th 2013, through close collaboration between SAP, HP and Microsoft, a **new world record** was achieved and certified by SAP for a three-tier SAP Sales and Distribution (SD) standard application benchmark, running on a set of 2-processor physical servers.

The application benchmark resulted in 42,400 SAP SD benchmark users, 231,580 SAPS, and a response time of 0.99 seconds, showcasing phenomenal performance using a DBMS server with just 2 physical processors of 16 cores and 32 CPU threads.

Not only was SAP ERP 6.0 (with Enhancement Package 5) running on SQL Server 2012, on Windows Server 2012 Datacenter, but **the configuration was completely virtualized on Hyper-V**. In addition, this is the first SAP benchmark with virtual machines configured with 32 virtual processors, and subsequently, the first with SQL Server running in a 32-way virtual machine. The result is also more than **30% higher** than a previous 2-processor/12 cores/24 CPU threads, virtualized configuration running on **VMware vSphere 5.0**.

It's clear from this benchmark that with the massive scalability and enterprise features in Windows Server 2012 Hyper-V, along with HP's ProLiant BL460c Gen8 servers, 3PAR StoreServ Storage and Virtual Connect networking capabilities, customers can virtualize their mission critical, tier-1 SAP ERP solution with confidence.

You can find the full details of the benchmark on the [SAP Benchmark Site](#), and you can also read more information about running SAP on Windows Server, Hyper-V & SQL Server, over on the [SAP on SQL Server Blog](#).

For more details visit: <http://www.sap.com/benchmark>

Note: - Benchmark performed in Houston, TX, USA on June 8, 2013. Results achieved 42,400 SAP Standard SD benchmark users, 231,580 SAPS and a response time of 0.99 seconds in a SAP three-tier configuration SAP EHP 5 for SAP ERP 6.0. Servers used for Application servers: 12 x ProLiant BL460c Gen8 with Intel Xeon E5-2680 @ 2.70GHz (2 processors/16 cores/32 threads) and 256GB using Microsoft Windows Server 2012 Datacenter on Windows Server 2012 Hyper-V. DBMS Server: 1 x ProLiant BL460c Gen8 with Intel Xeon E5-2680 @ 2.70GHz (2 processors/16 cores/32 threads) and 256GB using Microsoft Windows Server 2012 Datacenter on Windows Server 2012 Hyper-V using Microsoft SQL Server 2012 Enterprise Edition

VMware ESXi 5.0 based benchmark performed in Houston, TX, USA on October 11, 2011. Results achieved 32,125 SAP Standard SD benchmark users, 175,320 SAPS and a response time of 0.99 seconds in a SAP three-tier configuration SAP EHP 4 for SAP ERP 6.0. Servers used for Application servers: 10 x ProLiant BL460c G7 with Intel Xeon X5675 @ 3.06GHz (2 processors/12 cores/24 threads) and 96 GB using Microsoft Windows Server 2008 Enterprise on VMWare ESX 5.0. DBMS Server: 1 x ProLiant BL460c G7 with Intel Xeon X5675 @ 3.06GHz (2 processors/12 cores/24 threads) and 96 GB using Microsoft Windows Server 2008 Enterprise on VMWare ESX 5.0 using Microsoft SQL Server 2008 Enterprise Edition

How does VMware Compare?

With Windows Server 2012, and subsequently, Windows Server 2012 R2, Microsoft is bringing to market a virtualization platform that is highly optimized, offering the highest levels of performance for key workloads at a price point significantly lower than that of VMware, and in addition, Microsoft strives to provide customers with the confidence that modern, mission-critical tier-1 workloads run extremely well on the platform. Through collaboration with partners, and engagement with 3rd parties for testing, we showcase our performance through a number of reports that reflect modern versions of the applications, running on modern versions of the platform.

VMware however, take a different approach, providing customers with performance test results validated only on legacy versions of the applications, such as the performance testing of Exchange 2007 on ESX 3.5, Exchange 2010 on vSphere 4.0, SQL Server 2008 on vSphere 4.0, and no validated testing results for SharePoint performance.

VMware is very forthcoming with positioning its performance characteristics of vSphere, but struggles to back this up with validated, published performance test results for current generation workloads on their most modern platform versions, leaving customers with a lack of confidence to run their newest tier-1 workloads on the 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.

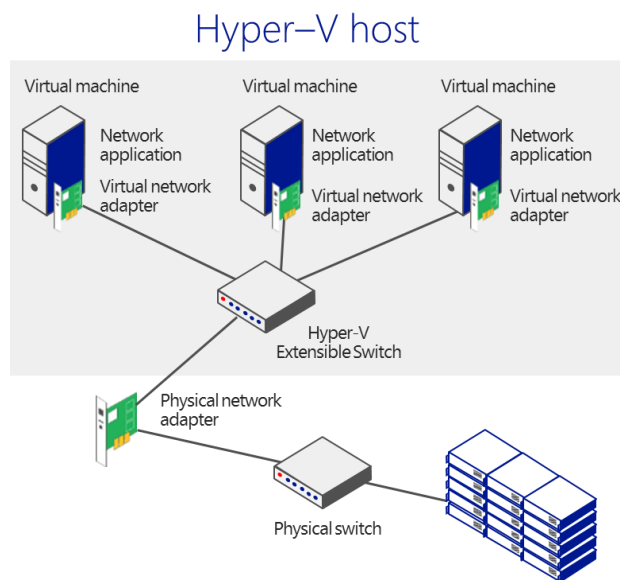


Figure 10 – Network with Hyper-V Extensible Switch

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.

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 those 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.

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.
- **5nine – Security Manager** provides both 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.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Extensible Network Switch	Yes	No	Replaceable
Confirmed Partner Extensions	5	No	2
Private Virtual LAN (PVLAN)	Yes	No	Yes
ARP/ND Spoofing Protection	Yes	No	vCloud/Partner
DHCP Snooping/DHCP Guard	Yes	No	vCloud/Partner
Virtual Port ACLs	Yes	No	vCloud/Partner
Trunk Mode to Virtual Machines	Yes	No	Yes
Port Monitoring	Yes	Per Port Group	Yes
Port Mirroring	Yes	Per Port Group	Yes

Whilst VMware offer an advanced distributed network switch, unfortunately, it is only available in the Enterprise Plus edition of vSphere 5.5, thus customers wishing to take advantage of the increased granularity, management capability and control, have to upgrade to the highest edition, at substantial cost. The VMware vSphere Hypervisor, unfortunately doesn't provide this capability. A key point to note however, is that the vSphere vSwitch, isn't open and extensible, but instead, closed and replaceable. Up until recently, Cisco were the only vendor to offer an alternative to the VMware vSphere Distributed Switch. IBM have released an alternative to this, however with Windows Server 2012 R2 Hyper-V, there are already multiple partners, including Cisco, NEC, 5nine and InMon, that have delivered extended functionality across a variety of different extension types, from packet inspection and filtering through to forwarding and intrusion detection, offering customers a greater set of choice for their specific environment. It's also important to note that so far, the approach from VMware's Partners has been more about replacement than integration, with the Cisco Nexus 1000V and the IBM System Networking Distributed Virtual Switch 5000V both effectively replacing the in-box vSphere Distributed Switch. Also, with System Center 2012 R2 Virtual Machine Manager, customers will be able to centrally control and manage the Hyper-V extensible switches across all hosts, providing greater control, and simplified management for larger environments.

Many of the more advanced networking capabilities within Windows Server 2012 R2 Hyper-V are unfortunately not present within the free vSphere Hypervisor, and even with vSphere, key security protection capabilities such as ARP and ND Spoofing Protection, DHCP Snooping Protection and DHCP Guard, along with Virtual Port Access Control Lists are only available through the purchase of additional technologies on top of vSphere 5.5, or within the network switch technologies from vendors such as Cisco. From a VMware perspective, prior to the release of vSphere 5.5, it was possible to purchase, at a per-VM cost, the vCloud Networking and Security product, which contained some of the more advanced networking capabilities found in-box with Hyper-V, however, since the

launch of vSphere 5.5, effective September 30th 2013, this [particular product line has been rolled up into the more expensive vCloud Suite](#).

This means that again, customers have to add additional, costly technologies in order to provide protection from these threats.

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 the IT Admin shape network traffic and enforce multitenant security in the data center. Unfortunately, this feature isn't available in vSphere Hypervisor but is available in the vSphere Distributed Switch, which is part of the vSphere Enterprise Plus edition, meaning customers have to upgrade to a higher edition in order to take advantage of the additional level of configuration granularity.

Finally, the Hyper-V Extensible Switch provides organizations with the ability to not only monitor individual ports within a vSwitch, but also mirror the traffic that is passing, to an alternative location for further analysis. With VMware vSphere Hypervisor however, all traffic on a Port Group or vSwitch, on which 'Promiscuous Mode' is enabled, is exposed, posing a potential risk to the security of that network. This lack of granularity restricts its usage in real world environments, and means that customers who require this level of protection have to upgrade to vSphere 5.5 Enterprise Plus, which has the Distributed Switch technology to provide the capability through features such as NetFlow and Port Mirroring.

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?

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.

How does VMware compare?

VMware has no capability within the vSphere Hypervisor or vSphere 5.5 that can enable the encryption of either VMFS, or the VMDK files themselves, and instead rely on hardware-based or in-guest alternatives, which add cost, management overhead, and additional resource usage.

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.

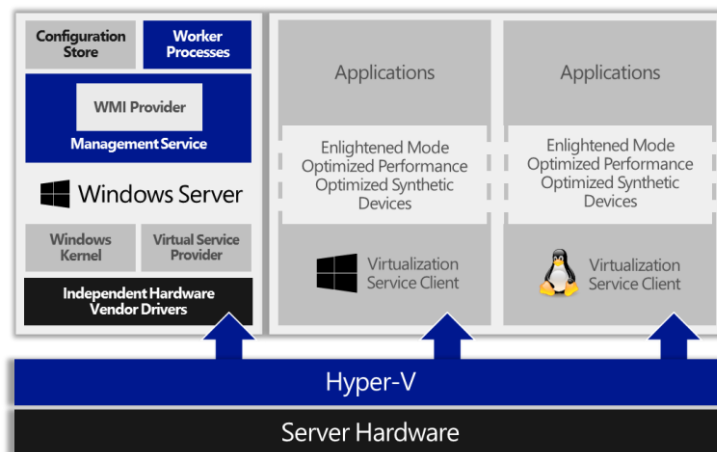


Figure 11 – Architectural view of Linux running enlightened on Hyper-V

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** – increased scalability and performance for Linux-based virtual machines along with better integration between Hyper-V services and the guest operating system.
- **Synthetic Frame Buffer Driver** – significantly improved graphics support for Linux-based virtual machines, including support for full HD resolutions (1920 x 1080)
- **Dynamic Memory Support** – Administrators can realize a 30-40% increase in server capacity when running Linux machines configured with dynamic memory. Dynamic Memory is also provided for Linux guests on older versions of Hyper-V, however administrators will realize greater levels of Dynamic Memory performance with more modern Linux distributions running on Windows Server 2012 R2 Hyper-V.
- **Live Virtual Machine Backup** - file-system snapshot driver is now available for Linux guests running on Hyper-V provides file-consistent snapshots without disruption to the running workload.

- **Online Resize of VHDX** – Allows administrators to flexibly allocate, or remove storage from a running virtual machine that contains a Linux guest OS.
- **Linux kdump/kexec Support** - Allows seamless creation of crash dumps using enlightened storage and network drivers and therefore no special configurations are required.
- **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** - 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.

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 12 – 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. 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.

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.

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.

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).

New in 2012 R2 - 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.

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.**

New in 2012 R2 - 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.

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.

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.

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.

New in 2012 R2 - 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

Hyper-V Network Virtualization

Isolating virtual machines of different departments or customers can be a challenge on a shared network. When these departments or customers must isolate entire networks of virtual machines, the challenge becomes even greater. Traditionally, VLANs are used to isolate networks, but VLANs are very complex to manage on a large scale. The following are the primary drawbacks of VLANs:

- Cumbersome reconfiguration of production switches is required whenever virtual machines or isolation boundaries must be moved, and the frequent reconfiguration of the physical network to add or modify VLANs increases the risk of an inadvertent outage.
- VLANs have limited scalability because typical switches support no more than 1,000 VLAN IDs (with a maximum of 4,095).
- VLANs cannot span multiple subnets, which limits the number of nodes in a single VLAN and restricts the placement of virtual machines based on physical location.

In addition to the drawbacks of VLANs, virtual machine IP address assignment presents other key issues when organizations move to the cloud:

- Required renumbering of service workloads.
- Policies that are tied to IP addresses.
- Physical locations that determine virtual machine IP addresses.
- Topological dependency of virtual machine deployment and traffic isolation.

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 R2 in one network subnet to servers running Windows Server 2012 R2 in a different network 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

New in 2012 R2 - Network Virtualization Routing | Outside the Virtual Network

Most customer deployments will require communication from the Hyper-V Network Virtualization (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.

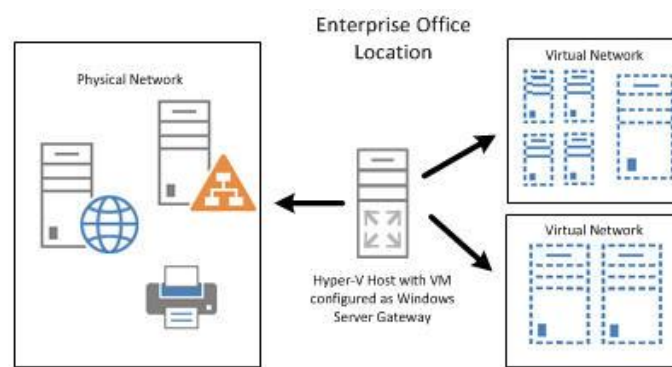


Figure 13 – Single location, with the WSG being used to connect VMs with physical infrastructure

The Windows Server Gateway (WSG), built into 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. This is included in the box, at no additional cost, allowing customers to fully embrace Network Virtualization in a scalable, production deployment, yet not be required to purchase costly, additional technologies.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
VM Live Migration	Yes	No	Yes
VM Live Migration with Compression	Yes	N/A	No
VM Live Migration over SMB/RDMA	Yes	N/A	No
1 GbE Simultaneous Live Migrations	Unlimited	N/A	4
10 GbE Simultaneous Live Migrations	Unlimited	N/A	8
Live Storage Migration	Yes	No	Yes
Shared-Nothing Live Migration	Yes	No	Yes
Live Migration Upgrades	Yes	N/A	Yes
VM Live Cloning	Yes	No	Yes
Network Virtualization	Yes	No	VXLAN/NSX
Network Virtualization Gateway	Yes	No	vCloud Suite
Flexible Linux Guest OS Support	Yes	Yes	Yes

As shown in the table, the flexibility and agility provided by the in-box features of Windows Server 2012 R2 Hyper-V are simply unmatched by VMware. The VMware vSphere Hypervisor supports none of the capabilities required for an agile infrastructure today, meaning customers have to purchase a more expensive vSphere 5.5 edition.

vSphere 5.5 Essentials Plus edition, and higher, now support vMotion (virtual machine live migration) yet on 1GigE networks, VMware restrict the number of simultaneous vMotions to 4, and on 10GigE, to 8. With Windows Server 2012 R2 Hyper-V, Microsoft supports an unlimited number of simultaneous live migrations, within the confines of what the networking hardware will support, with the process efficiently utilizing 100% of the available, dedicated live migration network to complete the process as quickly and efficiently as possible, with no interruption to the running virtual machines.

With Windows Server 2012 R2, Microsoft continues to push the boundaries with VM Live Migration, implementing new ways to utilize hardware investments, across both CPU, with Live Migration Compression, or with the network adaptors themselves, with Live Migration over SMB. Aside from supporting multiple network adaptors for vMotion in vSphere 5.0, VMware have offered little since then in terms of optimizing live migration performance. VMware still, with vSphere 5.5, do not support vMotion over RDMA-capable network hardware, nor have anything that can match Live Migration with Compression offered in-box with Hyper-V in Windows Server 2012 R2.

Just like virtual machine vMotion, Storage vMotion is unavailable in VMware vSphere Hypervisor, and is restricted to the Standard, Enterprise and Enterprise Plus editions of vSphere 5.5, available at considerable cost. In vSphere 5.1, VMware also introduced a feature, known in certain documentation, as Enhanced vMotion, which

enables the migration of a virtual machine between 2 hosts without shared storage. This feature was already available in all editions of Hyper-V, in the form of Shared-Nothing Live Migration.

When it comes to cloning virtual machines, without downtime, vSphere 5.5 has the ability to clone flexibly, however customers must use vCenter for this, meaning this capability is not available for customers just using the free vSphere Hypervisor. This is very different from the Microsoft approach, which provides customers with the ability to flexibly clone virtual machines using Hyper-V Manager GUI, or PowerShell, with either Windows Server 2012 R2 Hyper-V, or the free Hyper-V Server 2012 R2.

Finally, when it comes to Network Virtualization, it's clear that there are a number of significant benefits that can be realized through its implementation within enterprises and service providers, however, VMware's story around network virtualization is one of complexity and confusion, and ultimately for the customer, one of increased costs.

VMware initially led their networking virtualization push, back in 2011, with VXLAN, a collaboration between VMware, Cisco, and other industry leaders, to enable multi-tenant networking at scale, within the virtualized environment. When it finally released to market, VXLAN was packaged exclusively within the vCloud Networking and Security SKU, which, as with many of VMware's other features, was an expensive add-on to vSphere, and due to its requirement for the vSphere Distributed Switch, VXLAN required the Enterprise Plus edition of vSphere, meaning additional expense. Fast forward to July 2012, and VMware announce their intention to acquire Nicira, a business focused exclusively on software defined networking. This leaves customers with a very tricky decision – do they invest resources, time, effort, and funds, down a VXLAN route, knowing that Nicira was on the horizon, or do they wait, to see how the integration will unfold? Alternatively, do they embrace a clear strategy, clear from the start, with Hyper-V Network Virtualization, which, unlike VMware's features, is built into Windows Server 2012 R2 Hyper-V, at no additional cost.

Fast forward now to 2013, and with the release of vSphere 5.5, VMware NSX is VMware's networking virtualization platform, encompassing the Nicira capabilities, yet at time of publication, no pricing or packaging information is available, and a release date unconfirmed but it is unlikely to feature as part of the core vSphere SKU, with an addition to the vCloud Suite, or a standalone SKU more likely.

With Hyper-V Network Virtualization, core to the engine of Hyper-V, and managed at scale with System Center, customers can start to realize the benefits of Network Virtualization today, with a solution that has clear direction for the future.

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

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.

Windows Server 2012 R2 has a number of teaming and traffic distribution algorithms that ensure optimized performance and reliability, however new in Windows Server 2012 R2 specifically, is the Dynamic Mode algorithm. **Dynamic Mode** uses the concept of flowlets to achieve dynamic Load Balancing Fail Over (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.

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.

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.

As part of the Windows Server 2012 and Windows Server 2012 R2 releases, there have been a number of enhancements within Failover Clustering to ensure that it provides the most optimized, resilient platform for the key workloads virtualized on the cluster.

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).

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.

Affinity

With Windows Server 2012 and Windows Server 2012 R2 Failover Clusters, the administrator can define preferred and possible owners. Preferred ownership defines preference of VM location upon failover, and possible ownership, is where, for a given VM, you can configure the nodes which the VM has the possibility of failing over to.

Anti-Affinity

Allows the administrator to define groups of related VMs that they wish to keep apart, running on different physical cluster nodes. 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

Cluster-Aware Updating (CAU) is a key capability, built into Windows Server 2012 R2 that allows the IT admin 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, CAU automatically Live Migrates virtual machines from the physical nodes, before patching, ensuring that key virtualized workloads remain as continuously available as possible.

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 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:

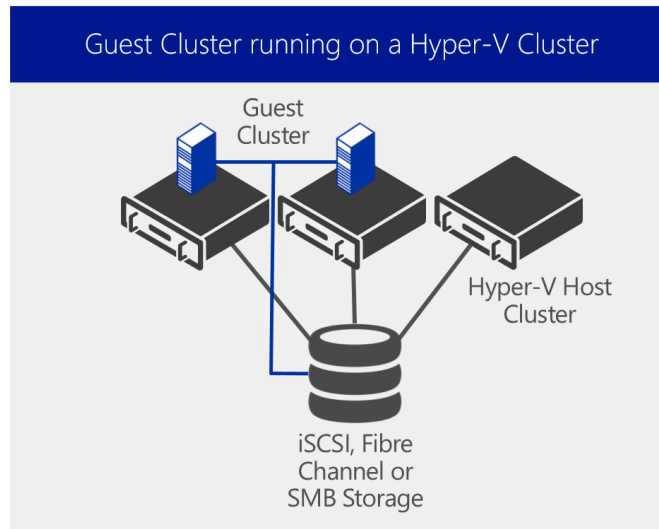


Figure 14 – A Guest Cluster running on a Hyper-V Cluster

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 configurations, 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 15 – 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.

New in 2012 R2 - 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.)

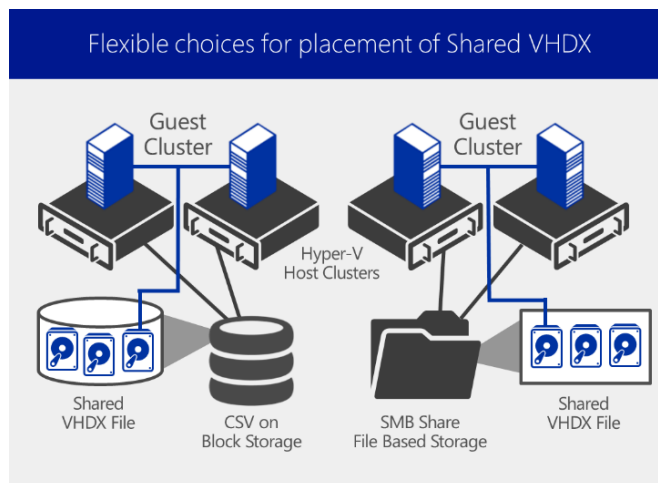


Figure 16 – A Guest Cluster using Shared VHDX on CSV/SMB Storage

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.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
NIC Teaming	Yes	Yes	Yes
Integrated High Availability	Yes	No	Yes
Nodes per Cluster	64	N/A	32
VMs per Cluster	8,000	N/A	4,000
Max Guest Cluster Size (iSCSI)	64 Nodes	5	5
Max Guest Cluster Size (Fiber)	64 Nodes	5	5
Max Guest Cluster Size (File Based)	64 Nodes	5	5
Guest Cluster with Shared Virtual Disk	Yes	Yes	Yes
Guest Clustering with Live Migration	Yes	N/A	No
Guest Clustering with Dynamic Memory	Yes	No	No
Guest OS Application Monitoring	Yes	N/A	Yes
Cluster-Aware Updating	Yes	N/A	Yes
Failover Priority, Affinity & Anti-Affinity	Yes	N/A	Yes

The table above shows that when it comes to comparing the clustering and high availability capabilities of Windows Server 2012 R2 Hyper-V and the vSphere Hypervisor, the restrictions placed on VMware's free edition become quickly evident. Whilst the vSphere Hypervisor does support integrated NIC Teaming for network card resilience, it is lacking many of the other, important resiliency features, meaning if customers were to virtualize important workloads on the platform, they would have to upgrade to a more expensive edition in order to provide some form of resiliency and protection for the workloads in question.

Windows Server 2012 R2 Hyper-V on the other hand, offers a number of resiliency and high availability features in the box. Integrated Failover Clustering provides the foundation for virtual machine resiliency upon host, and virtual machine failure, and extends the native protection into the guest operating system, ensuring that if application services start to exhibit problems, corrective action can be taken.

With vSphere 5.5, VMware introduced a new capability, known as vSphere App HA. vSphere App HA, similarly to Windows Server 2012 R2 VM Monitoring, looks inside the VM, at important service, and should a service fail, vSphere App HA policies will trigger the restart of a service, restart of the VM, or a reset if necessary. One of the key differences between this and the Microsoft implementation, is the additional infrastructure requirements for the VMware solution, with App HA requiring the provisioning of 2 additional virtual appliances to run the feature and in addition, agents to be installed in the protected VMs. Compare this with Microsoft's approach, which

requires no additional agents inside the VM, and the feature runs as a core host service, so no additional appliance deployment required, saving valuable host resources and administrative time and effort.

Windows Server 2012 R2 Hyper-V offers double the number of nodes in an individual cluster, when compared with vSphere 5.5, and scales the number of virtual machines within an individual cluster to 8,000, again, double that of vSphere 5.5. This provides large enterprises, and service providers with unprecedented scale to run significant numbers of workloads and achieve a significantly stronger return on investments.

Customers who have embraced the standalone vSphere Hypervisor don't have the ability to construct resilient clustered infrastructures, unless they upgrade to a costly edition of vSphere 5.5, however customers wishing to construct virtual machine guest clusters can use the standalone vSphere Hypervisor, or alternatively, vSphere 5.5. VMware's support for guest clusters is severely lacking in comparison with Microsoft's flexible offerings. Customers who have invested in iSCSI storage can create guest clusters on the vSphere Hypervisor or on vSphere 5.5, using iSCSI connectivity, however with vSphere 5.5, VMware support clusters running Windows Server 2008 R2 or newer, and [restrict customers to just 5 node guest clusters](#). The same scalability limits are imposed for guest clusters utilizing virtual fibre channel, or guest clusters running over network file storage. Compare this with Windows Server 2012 R2 Hyper-V, which, for a Windows Server 2012 or 2012 R2 guest cluster, supports up to 64 nodes, over iSCSI, Virtual Fiber Channel, or SMB 3.0, for complete flexibility and unmatched scale.

Also, with the new Shared VHDX support in Windows Server 2012 R2 Hyper-V, customers have yet another flexible option for storage when deploying guest clusters, without needing to expose the underlying storage through to virtual machines. VMware too, support guest clusters with shared virtual disks, yet the problem with the VMware implementation, is that the whole cluster has to reside on the same physical ESXi host, meaning no resiliency for host failures, and the whole cluster would go down.

It's important to note that whilst Windows Server 2012 R2 Hyper-V clearly provides a significantly more comprehensive guest clustering capability than VMware in terms of storage integration and support, it also doesn't require customers to sacrifice other features and functionality to work effectively. A virtualized guest cluster on Windows Server 2012 R2 Hyper-V supports features such as virtual machine Live Migration, for flexibility and agility, and Dynamic Memory, to ensure the highest levels of density. Compare this with VMware, who, whilst restricting customers to a maximum of 5 nodes, they also restrict customers from migrating the guest cluster nodes using vMotion, migrating disks with Storage vMotion, and additionally, direct customers to disable memory overcommit on those guest cluster nodes, sacrificing density. These are just a few of the many limitations with VMware vSphere guest clustering.

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, ensuring optimal use of disk space.

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 integrates directly with Windows Server Backup and then performs block-level incremental backup. To increase security, the data is compressed and encrypted before leaving the server.

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.

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 or 15 minutes.

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. This means customers can replicate a VM from Site A -> Site B -> Site C, providing even greater protection against potential disasters.

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 17 – 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.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
Incremental Backup	Yes	No	Yes
Inbox Software-Based VM Replication	Yes	No	Yes

As you can see from the table above, both Microsoft and VMware offer backup capabilities to protect the important virtualized workloads, yet VMware’s backup capabilities, found within the vSphere Data Protection product, are not available for the free vSphere Hypervisor, which is a very different approach to that of the free Hyper-V Server 2012 R2, which does include Windows Server Backup, to enable organizations to protect their key workloads, and very low cost. From a VM Replication perspective, it’s a similar story, with customers having to purchase a vSphere SKU to take advantage of the vSphere Replication that is available with vSphere, however, as the next table will show, even customers who choose to purchase vSphere will see that the replication features offered by VMware and Microsoft differ considerably.

Replication Capability	Hyper-V Replica	vSphere Replication
Architecture	Inbox with Hypervisor	Virtual Appliance
Replication Type	Asynchronous	Asynchronous
Recovery Time Objective (RTO)	30s, 5m, 15m	15m – 24h
Replication	Tertiary	Secondary
Planned Failover	Yes	No
Unplanned Failover	Yes	Yes
Test Failover	Yes	No
Simple Failback Process	Yes	No
Automatic Re-IP Address	Yes	No
Point in Time Recovery	Yes – 24 points	Yes - 24 points
Orchestration	PowerShell, SC Orchestrator, HVRM	No Scripting, SRM only

As you can see from the table, the capabilities offered in the box with Hyper-V Replica are very different from those offered by VMware with vSphere Replication. Firstly from an architectural perspective, Hyper-V places no reliance on additional appliances in order to manage the replication, freeing up valuable resources and eliminating points of failure for replication.

Also, Hyper-V Replica offers near-synchronous replication, down to every 30 seconds, meaning it's significantly more useful for mission-critical workloads, than that offered by vSphere Replication, which can only replicate data as low as every 15 minutes, and only between 2 locations. Contrast that with Hyper-V Replica, and you'll find a VM that can be replicated from Site A, on to Site B, and from B, on to Site C, providing a significantly differentiated protection capability against multi-site loss.

Hyper-V Replica also offers more granularity around performing a planned or unplanned failover, testing a failover, and configuring VM NIC settings in advance, before failover. VMware's vSphere Replication only offers the administrator to perform an unplanned failover i.e. after a disaster has taken place, with no ability to test in advance, or fail over cleanly when pre-warned of a disaster. Customers must upgrade, at considerable cost, to Site Recovery Manager, to unlock some of this core functionality. Core functionality that is included with Hyper-V Replica in the box. SRM also provides vSphere Replication with an automation engine, which without it, there would be no way to automate vSphere Replication. Hyper-V Replica on the other hand, can be automated via PowerShell, at extreme low cost, or can utilize investments in System Center Orchestrator, or the Windows Azure Hyper-V Recovery Manager, for orchestrated failover between sites.

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 a couple 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.

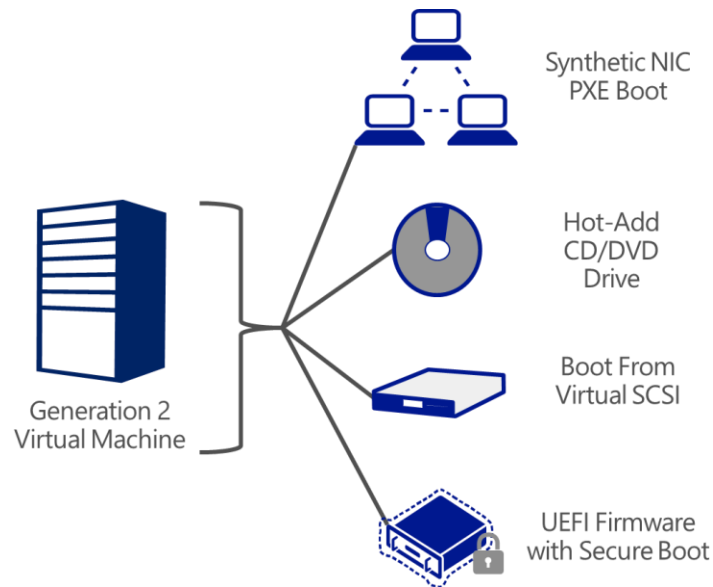


Figure 18 – Key Changes with Generation 2 VMs

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.

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 19 – 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.

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.

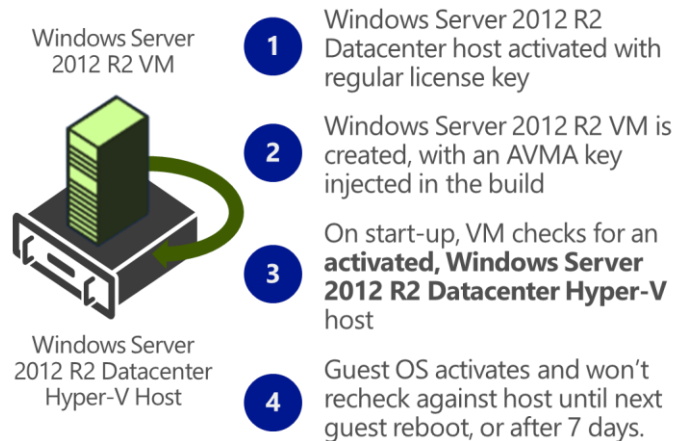


Figure 20 – Automatic VM Activation

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.

How does VMware compare?

Capability	Windows Server 2012 R2 Hyper-V	VMware vSphere Hypervisor	VMware vSphere 5.5 Enterprise Plus
VMs with Secure Boot & UEFI Firmware	Yes	No	No
Enhanced VM Administration Experience	Yes	No	No
Automatic VM Activation	Yes	No	No

Generation 2 virtual machines bring some very significant benefits. Aside from reducing the number of emulated, or legacy devices, and bringing faster boot and OS install times, Generation 2 virtual machines provide

the VM with UEFI firmware. UEFI firmware brings a number of advantages, with one of the most significant being Secure Boot, which will prevent potential rootkits that could be injected into the pre-boot-up processes of an operating system. Unfortunately for VMware, they have no similar protection capabilities, and rely solely on the more legacy BIOS approach for the virtual machines.

When it comes to the Enhanced Session Mode, Hyper-V offers administrators not only a much richer, more streamlined management experience for administering and interacting with the VMs, than compared with the vCenter experience, it's what Enhanced Session Mode enables from a troubleshooting perspective which is perhaps most important. The ability of Hyper-V to provide RDP over VMBus, meaning administrators can still RDP into the VM when VM network connectivity is down, is very valuable from a troubleshooting perspective. VMware can offer nothing similar, even in vSphere 5.5.

Finally, with the new Automatic VM Activation, VM activation for customers of all shapes and sizes has been simplified considerably. Through the integration between activated host, the Hyper-V hypervisor, and the VM Guest OS, the VMs can be safely and securely activated using a generic key, without needing to administer licensing and product key management manually, or through centralized management tools. With VMware however, as there isn't the level of host integration, it means customers still have to rely on traditional mechanisms for license activation.

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 4 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 all editions. 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, which, unlike VMware, are simply built into Hyper-V, which enables customers to securely and easily isolate and control access to their key workloads inside the virtualized environment, at significantly lower cost.

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.

In addition, we've also looked at 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. These are just a few of the innovative examples that push Hyper-V beyond the capabilities of VMware vSphere.

For more information on Windows Server 2012 R2 Hyper-V, visit <http://www.microsoft.com/en-us/server-cloud/solutions/virtualization.aspx>