solarwinds

Hyper-V® 2012 vs. vSphere™ 5.1

Understanding the Differences

Scott D. Lowe

Virtualization Blogger & vExpert, Founder & Consultant of The 1610 Group

Table of Contents

Introduction

For years, there has been a feature battle between VMware® vSphere™ and Microsoft® Hyper-V®, a battle which has heated up with both companies releasing major new editions of their hypervisors in 2012. In previous editions, vSphere has handily won the side-by-side feature battle, although Hyper-V was certainly no slouch; the product was able to hold its own and did so for many organizations that made the decision that it's the right solution for their needs.

With the release of Hyper-V 2012 and vSphere 5.1, however, Microsoft is now able to do much more than just "hold its own." In fact, today's comparison, which is the topic of this paper, shows a Microsoft that has reached feature parity with vSphere in many ways and even surpassed VMware in others.

vSphere administrators may not have the opportunity or desire to really put Hyper-V through its paces, particularly given vSphere's continued leadership position in the market. In order to help vSphere administrators make the knowledge leap between the two platforms, this paper will outline the differences between the currently shipping products from both VMware and Microsoft: vSphere 5.1 and Hyper-V 2012. In this paper, you will be introduced to key features in the products and be exposed to the feature deltas that may exist.

Product-Level Overview

Both VMware and Microsoft offer their products in multiple editions depending on the needs of the customer and both companies provide a free edition of their respective hypervisor tools for customers that do not have advanced needs.

Supported Guest Operating Systems

In breadth of operating system support, VMware's vSphere has been and remains the clear winner, although Microsoft has made steady strides in Hyper-V. For many, however, Hyper-V's support for the world's most common operating systems is enough.

Hyper-V supports all Microsoft operating systems – both client and server – since Windows® 2003 SP2 and Windows XP SP3. In addition to supporting Windows operating systems, Hyper-V supports:

- CentOS 5.2-5.8 and CentOS 6.0 and 6.3
- Red Hat® Enterprise Linux® 5.7 – 5.9, 6.0 –6.3
- SUSE® Linux Enterprise Server 11 with Service Pack 2
- openSUSE® 12.1
- Ubuntu® 12.04

With Hyper-V 2012, Microsoft has dropped support for versions of Windows prior to Windows Server 2003 as well as support for older editions of various Linux distributions.

VMware, on the other hand, provides a much more comprehensive level of support for guest operating systems. The company supports Microsoft products all the way back to MS DOS 6.22 and Windows 3.1, although these really old products have support that is deprecated in vSphere 5.1. On the Linux front, VMware supports products from just about every distribution out there, including CentOS, Mandrake, Novell and Red Hat, among others. In addition to supporting most Linux variants, vSphere also provides support for Mac OS® X 10.7 and other versions of OS X Server, FreeBSD®, OS/2® Warp, NetWare®, and Solaris®. For a complete list of supported operating systems, refer to VMware's **Guest Operating System Installation Guide**.

Scalability Comparison on a Per-Edition Basis

In previous comparisons of vSphere and Hyper-V, the products differ significantly in their ability to scale to serve large organizations with vSphere generally winning in a side-by-side comparison. Today, in a number of different ways, Hyper-V has closed the gap and even enjoys front-runner status in some areas. The chart below details the scalability differences in the various editions of the two company's products. You will note that many of the metrics now favor Hyper-V. What's most interesting is the number of features that are supported in the free edition of Hyper-V even without running System Center 2012 Virtual Machine Manager.

	vSphere 5 (no vCenter)	vSphere 5 (w/ vCenter)			Hyper-V (no VMM)	Hyper-V (w/ SC VMM)		
Feature	Free	Std	Ent	Ent Plus	Free	Free	Std	DC
Max host processors	160	160	160	160	320	320	320	320
Max virtual SMP (guest)	8	8	32	64	64	64	64	64
Max host RAM (GB)	32	2048	2048	2048	4096	4096	4096	4096
Max RAM per VM (GB)	32	1024	1024	1024	1024	1024	1024	1024
Max guest disk size (TB)	2 64 (RDM)	2 64 (RDM)	2 64 (RDM)	2 64 (RDM)	2 (VHD) 64 (VHDX) 256+ (RAW)	2 (VHD) 64 (VHDX) 256+ (RAW)	2 (VHD) 64 (VHDX) 256+ (RAW)	2 (VHD) 64 (VHDX) 256+ (RAW)
Failover nodes	○	32	32	32	64	64	64	64
Memory overcommit/ dynamic mem.	●	●	●	●	●	●	●	●
Transparent page sharing	●	●	●	●	○	○	○	○
Live workload migration	○	●	●	●	●	●	●	●
Live storage migration	○	●	●	●	●	●	●	●
Max guests per host	512	512	512	512	1024	1024	1024	1024
Automatic workload migration	○	○	●	●	○	●	●	●
Automatic storage migration	○	○	○	●	○	●	●	●
Enhanced switching	○	○	○	●	●	●	●	●
Virtual instance rights (Windows)	○	○	○	○	○	○	2	Unlimited
I/O passthrough	VMDirectPath	VMDirectPath	VMDirectPath	SR-IOV + VMDirectPath	SR-IOV (works w/ migration)	SR-IOV (works w/ migration)	SR-IOV (works w/ migration)	SR-IOV (works w/ migration)
Hypervisor licensing model	N/A	per proc	per proc	per proc	N/A	N/A	per proc	per proc

○ **Feature not supported** ● **Feature supported**

Sources: Configuration Maximums (VMware), Requirements and Limits for Virtual Machines and Hyper-V in Windows Server® 2012

Technical Features Comparison

As mentioned, today's leading hypervisors have become very close to one another in terms of features and capabilities. Today, vSphere 5.1 and Hyper-V 2012 are relatively closely matched when it comes to breadth and depth of features. In fact, in some ways, Microsoft has an edge in many ways since many of Hyper-V 2012's best features are usable even with the free edition of Hyper-V. The free

edition of vSphere is a pretty crippled product that has an artificial RAM limit and can't be managed by vCenter™, thus eliminating potential use of many VMware features.

Of course, vSphere still reigns supreme in the market due to the product's rock solid performance and ease of use. While Microsoft has absolutely caught up to VMware in most areas, vSphere still tops the list when it comes to ease of use.

Architecture & Hypervisor Footprint Differences

Perhaps some of the most obvious differences between VMware and Hyper-V lie in the product architecture and footprint. Whereas a VMware vSphere Hypervisor installation requires only 144 MB of disk space, a Hyper-V 2012 installation requires a minimum of 5 GB for a Server Core installation and around 10 GB or so when a full Windows installation is selected.

The primary reason for this significant space difference lies in the underpinnings of the two hypervisor solutions. Hyper-V relies on the installation of the general purpose Windows Server 2012 operating system while vSphere does not have such weight attached to it.

On the architecture front, Hyper-V is often mistaken for what's known as a Type 2 hypervisor when, in reality, both it and vSphere are Type 1 hypervisors. With a Type 1 hypervisor, often referred to as a "bare metal" hypervisor, the virtualization software sits directly atop the hardware, managing access to said hardware. With Type 2 hypervisors, such as VirtualBox from Oracle® and VMware Workstation, the hypervisor is installed inside a host operating system and the hypervisor software simply operates like any other host-based program.

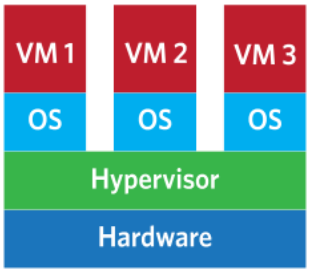

Figure 1: Type 1 hypervisor

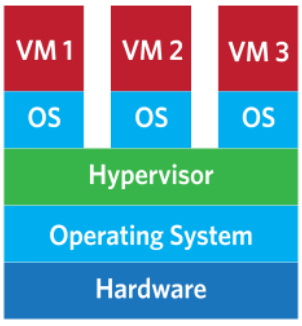

Figure 2: Type 2 hypervisor

When the Hyper-V role is enabled on a Windows Server 2012 machine, the existing Windows instance becomes a virtualized root partition. From an architecture standpoint, Hyper-V's reliance on and support of Windows as its base means that Hyper-V can be used on a much wider variety of hardware than vSphere, which is pretty restrictive when it comes to hardware. The Windows hardware compatibility list (HCL) provides much broader hardware support than vSphere's relatively narrow list of supported hardware.

If you dig a little deeper into the technical details, you will also find that the Hyper-V architecture is based on what's called a microkernelized hypervisor while vSphere is monolithic in nature. This microkernelized design lends flexibility and security to the hypervisor model by isolating the virtual machines from one another with little shared code, such as drivers. Under Hyper-V, more synthetic drivers are used, which can boost overall service performance. vSphere's architecture revolves around a more monolithic core which includes many shared drivers as well as the virtualization stack.

Frankly, the kernel type – microkernelized vs. monolithic – probably won't play much of a part in a hypervisor decision at present, but understanding that both vSphere and Hyper-V both leverage bare metal capabilities is critical, particularly since Hyper-V's implementation model is fraught with confusion.

CPU Contention Management

Regardless of the system running, the operating system always has to schedule CPU resources appropriate to the current assignment. As the system becomes charged with the support of additional disparate workloads, however, the issue of CPU contention – that is, the ability for the system to respond to processing requests – becomes an issue. However, resource contention is generally not a major issue at the low end of the virtualization spectrum. It begins to come into play as systems are pushed to their limits, which forces the operating system to more carefully dole out increasingly scarce resources.

vSphere CPU Scheduling Controls

Both vSphere and Hyper-V include manual methods for adjusting CPU scheduling between virtual machines. There are three settings for each virtual machine in vSphere that can be adjusted:

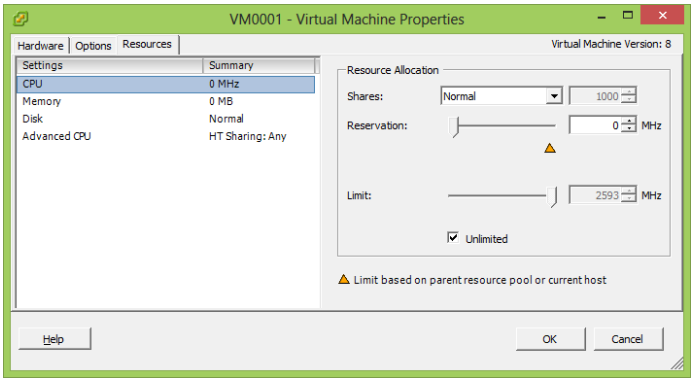

- **Shares**. Shares are used to dictate the relative performance of a virtual machine. If a virtual machine has a share value that is half of another, it's entitled to only half the CPU resources.
- **Reservation**. A reservation is a guarantee that a virtual machine will receive at least the resources that are specified in MHz.
- **Limit**. The value specified here limits the ability of the virtual machine to consume unlimited resources. This is useful in situations in which an individual virtual machine is consuming too much.

At the host level, vSphere has a powerful CPU scheduling mechanism in place that ensures that virtual machines receive attention from the system. VMware has produced a **white paper** that goes into great technical depth for how this scheduling is achieved.

Hyper-V CPU Scheduling Controls

Hyper-V works to address CPU contention issues in a couple of different ways. First, each virtual machine has manual CPU settings that the administrator can adjust to meet business needs.

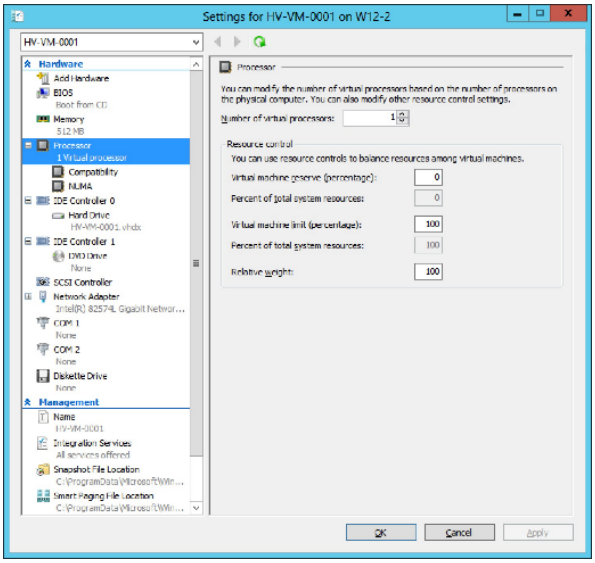

- **Virtual machine reserve** (percentage). Allows the reservation of a portion of the server's total processing resources for this virtual machine. Consider this scenario: This virtual machine is running a mission-critical workload. There must always be CPU resources available to serve this VM's workload. By default, a virtual machine is guaranteed/reserved 0% of the host's resources. An administrator can set this to a non-zero value to reserve resources.

- **Virtual machine limit** (percentage). On the flip side, an administrator can also limit how much of a host's processing resources can be consumed by a single virtual machine. This setting is useful for instances in which a virtual machine might be attempting to consume too many resources.
- **Relative weight**. If the two settings above are a bit too exacting, an administrator can take a different approach to determine how much processing power should be consumed by the virtual machine. The relative weight option allows the weighting of this virtual machine against others. By default, every virtual machine gets a weight of 100. If a VM should have lower priority, provide a lower number.

Shared Memory/Memory Management Differences

vSphere currently handily trumps Hyper-V in the area of memory management. The two sections below provide you with technical details about how each solution implements advanced memory management techniques. Hint: If memory management is your key technical driver, stop reading now and stick with vSphere.

VMware Memory Management Techniques

VMware has perfected some powerful techniques by which RAM can be managed and optimized on a vSphere host in order to provide additional scalability on a per-host basis and to keep a host operating a peak levels.

- **VMware Oversubscription/Overcommit**. VMware allows administrators to assign more aggregate RAM to virtual machines than is actually physically available in the server. This works because VMware actively monitors all virtual machines and can take RAM from virtual machines that aren't using their full allocation and assign that RAM to virtual machines that need it. Accomplished through a mechanism called the Idle Memory Tax, VMware progressively "taxes" virtual machines based on the amount of idle RAM they have.
- **Transparent Page Sharing**. When there is idle CPU time, vSphere looks for pages located across virtual machines that can be matched with one another and shared. This is basically a de-

duplication method applied to RAM rather than storage. For organizations that tend to use the same operating system for many virtual machines, the memory impact can be substantial.

- **Guest Ballooning**. When a particular virtual machine needs to increase its memory usage, it can request memory from other virtual machines. When VMware Tools is installed inside a virtual machine, along with everything else is a memory balloon process. The guest operating system can swap processes out to help free up memory that is then assigned to the balloon. When other VMs request memory, it can be lent via this balloon.
- **Memory compression**. Memory compression is a technique that is used to prevent the hypervisor from needing to swap memory pages to disk when RAM becomes limited. Memory compression attempts to fit multiple pages of RAM into a smaller number of pages in order to postpone for as long as possible the need for the hypervisor to swap to disk. Disk swapping is expensive in terms of performance.

Hyper-V Dynamic Memory

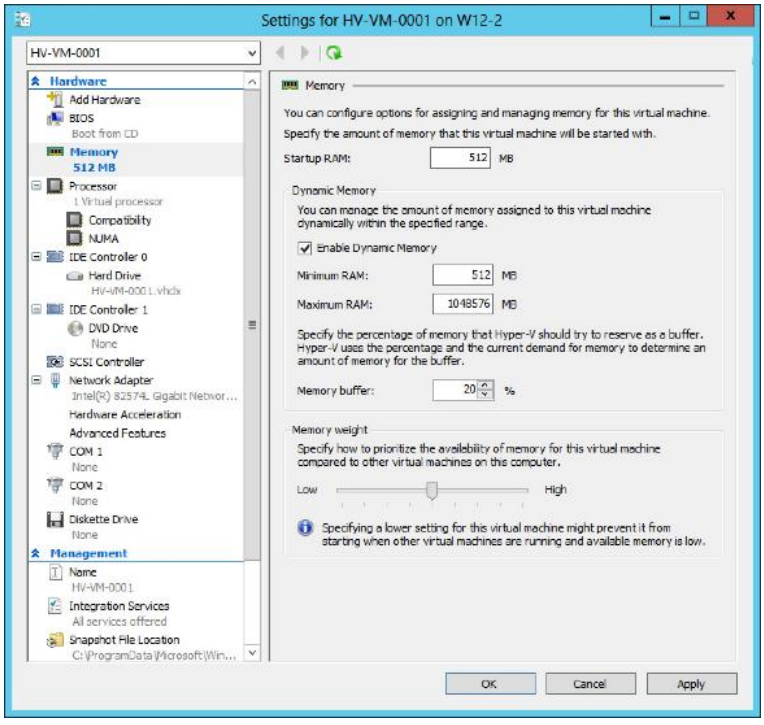

Whereas VMware has a number of different techniques to manage and optimize RAM, Hyper-V has just two.

The first method is known as **Dynamic Memory**, and was originally introduced in Windows Server 2008 R2 SP1. Dynamic Memory relies primarily on a process similar to the Guest Ballooning feature described previously.

Included in the Hyper-V Integration Services that are installed inside a virtual machine is a driver called the Dynamic Memory Virtual Service Consumer (DM VSC). The DM VSC is responsible for monitoring the memory usage of the guest operating and tracking changing guest RAM needs. This information is reported to the host, which is responsible for deciding whether to give or take RAM from the virtual machine for other purposes.

It's easy to see how this feature could reduce a virtual machine's memory to dangerous levels, leaving it unable to respond to a sudden environmental change that creates the need for more memory. To prevent this situation, Hyper-V provides a (default) buffer of 20% of unused memory.

Hyper-V exposes the Dynamic Memory configuration options as a part of the virtual machine settings. An administrator provides both startup and maximum RAM variables along with the desired memory buffer. The administrator also has the opportunity to provide a memory weight value to the virtual machine, which allows the administrator to ensure that mission critical virtual machines are serviced before less critical needs.

Dynamic Memory in Hyper-V is much more exposed to the administrator than RAM management options in VMware. Although the solution doesn't provide the same "overcommit" capability as VMware, it does allow the administrator to attempt to increase virtual machine density through advanced RAM management.

The second method, **Smart Paging**, is a new feature in Hyper-V 2012. As administrators configure Hyper-V guest virtual machines to maximize consolidation ratios, memory can become a limiting factor. For many workloads, ongoing memory usage is lower than memory usage required at boot time. Smart Paging steps in when a virtual machine is being rebooted and there is no available physical memory and no memory can be reclaimed from other virtual machines running on the host. Smart Paging enables the temporary use of storage as a memory cache so that the virtual machine can be provided with enough RAM to get started. Once started, the workload can fall back to its normal operational needs and stop using relatively slow disk as a temporary paging file.

This feature enables administrators to maintain high levels of virtual machine density without having to worry about boot-time RAM needs that may arise. However, it's important to keep in mind that this feature works only when a virtual machine has been rebooted, not when it's being started from a powered off state.

Virtual Machine Storage Capabilities Compared

Storage is another critical shared resource in a virtual environment. Only with the right kind of storage can the hypervisor products achieve their true potential. However, even before the issue of

advanced features as they relate to storage can be discussed, administrators must understand the storage capabilities and limitations that exist in both vSphere and Hyper-V.

Storage Types

The chart below outlines the storage types that are supported by each product. Hyper-V has added support for a much broader array of storage technologies, particularly with the release of the Server Message Block 3 (SMB 3) protocol.

Technology	Description	vSphere	Hyper-V
DAS	Directly attached storage	•	✓
NAS	Network attached storage	•	✓
FC	Fibre Channel	•	✓
iSCSI	Internet SCSI	•	✓
FCoE	Fibre Channel over Ethernet	•	✓

Thin Provisioning

Thin provisioning is a method by which administrators can have the best of both worlds when it comes to storage. No one wants to allocate too little disk space for mission critical services. On the other hand, no one wants to waste expensive shared disk space on a SAN, either. Thin provisioning makes it possible for administrators to allocate the space they believe they may ultimately need for a service without actually having to dedicate the space right now.

Thin provisioning is supported in both Hyper-V and vSphere. In addition, both vSphere 5.1 and Hyper-V 2012 support free disk space reclamation methods associated with thin provisioning. vSphere 5.1 accomplishes this goal through the introduction of a new virtual disk type, called the Space Efficient Sparse Disk. In Hyper-V 2012, the underlying Windows Server 2012 has added support for the TRIM API, which flows through to Hyper-V virtual disks.

Linked Image Support

Thin provisioning is one method by which disk space can be conserved in a virtual environment. Another way to achieve disk capacity savings is to simply link images to one another. For example,

suppose you're running fifty Windows Server 2012 virtual machines. All fifty of these machines use the same base disk image, so why not just create the image once and then point another forty-nine virtual machines at the base image?

This is an area in which Microsoft remains well ahead of VMware in most scenarios. Although VMware provides a feature known as "linked clones" in VMware View and vCloud™ Director, vSphere does not yet provide this capability natively.

Hyper-V provides a feature that Microsoft calls Differencing Disks, but with a caveat. A differencing disk (called a child disk) is a virtual hard disk linked to a master image (called a parent disk). A differencing disk stores and isolates any changes that are made to that master image on a per-VM basis. Through the use of this feature, administrators can create a master base image and simply link other virtual machines to this image and save a lot of disk space.

But, as mentioned, there is a caveat: Microsoft has traditionally recommended against using the differencing disks feature be used in a production environment due to performance issues that are associated with these types of disks.

vSphere VMFS vs. Microsoft VHDX

One of VMware's claims to fame has been VMFS (Virtual Machine File System), the incredibly powerful and scalable file system that enables much of VMware's feature set. VMFS delivers a number of features and capabilities that are not found in Microsoft's competing VHD (Virtual Hard Disk) and newer VHDX formats. In a cluster environment, VMFS truly shines. VMFS allows the hot add and removal of hosts to and from the cluster without the need to interrupt other running workloads. Although it's possible to add new Hyper-V hosts to a cluster, it's still not as seamless as the vSphere experience.

Microsoft does provide Cluster Shared Volume (CSV) support in Hyper-V, but the implementation of this feature is far more complicated than VMware's clustering implementation.

Both VMware and Microsoft allow administrators to make direct use of disks through respective pass-through techniques. For vSphere, this feature is known as Raw Device Mapping (RDM). In the

Hyper-V world, these are known as "Pass-Through Disks." In this area, there isn't much difference. Pass-through techniques can be used to improve overall performance.

On the storage management front, Microsoft has made great strides with regard to storage management and has closed the gap on a number of features that were missing from previous editions of Hyper-V and Virtual Machine Manager:

- **Centralized management of datastores**. A single location in which all datastores can be managed in order to provide more visibility into the environment. VMM now displays all stores.
- **Storage Management Initiative Specification (SMI-S) support**. Standardized monitoring of storage. Hyper-V 2012 now supports SMI-S.

However, Microsoft still does not have an answer to vSphere's Storage DRS (SDRS) feature, which provides a way to automatically place VMs to load balance Storage I/O demands.

Power Management

One of the great promises of virtualization is the possibility of reducing costs by combining workloads onto fewer physical machines, which also reduces the amount of electricity consumed in aggregate. That's just the first possibility, though. With advanced workload migration techniques, virtualization offers the possibility for optimizing running workloads by consolidating them onto as few hosts as possible and moving the remaining hosts into a power-save mode. In this way, organizations can make further gains in virtualization-related cost savings.

VMware Distributed Power Management (DPM)

Although there is some question in the VMware community about the actual rewards that can be reaped from DPM, it works exactly as described above and, theoretically, has the potential to reduce power usage. DPM automates the process of energy conservation, leaving the administrator free to focus elsewhere. When power requirements dictate, hosts are brought online to service the burgeoning workload and taken offline when needs lessen. DPM is available only in the Enterprise and Enterprise Plus editions of vSphere 5.1.

Hyper-V Power Optimization

In Hyper-V 2012 and when coupled with the Virtual Machine Manager 2012 management platform, Microsoft has introduced a feature known as Power Optimization. As long as there are a sufficient number of Hyper-V hosts in a cluster to meet quorum requirements, Power Optimization can shut down hosts that are not necessary to meet cluster performance and resource requirements, thereby optimizing the power usage of the cluster. As resource requirements in the cluster increase, Power Optimization can also restart hosts in order to add resources.

Virtual Network Features Comparison

One major change that was wrought with the rise of virtualization was the need for previously disparate IT teams to work more closely than ever before. Nowhere is this truer than when it comes to networking. With virtualization on the rise, the systems and networking teams need to collaborate and cooperate much more than was necessary in the past. Each hypervisor brings with it different capabilities, as described below.

Networking Features

Both Hyper-V and vSphere provide full support for foundational networking elements including VLAN support and support for jumbo frames, which can boost iSCSI storage performance. Both also provide full support for the emerging IPv6 standard as well as for bandwidth, and reliability-enhancing NIC teaming features, which have been significantly improved in Hyper-V 2012.

Further, both products provide support for different types of offloading, which can reduce host processor usage and increase potential consolidation ratios. In previous editions of these products, vSphere was the clear winner in the features race. However, Microsoft has all but closed the feature gap in Hyper-V 2012, particularly with regard to cluster-wide switching, which Microsoft makes available even in their free product, but for which vSphere requires customers to invest in their Enterprise Plus SKU.

vSphere

- **TCP Segmentation Offload**. The process of encapsulating data packets into Ethernet-sized frames requires some processor overhead. With TCP Segmentation Offload, the host TCP/IP stack can submit frames of up to 64 KB in size to the NIC at which point the NIC then repackages these frames into sizes that fit inside the network's maximum transmission unit (MTU) size.
- **NetQueue**. Enables the system to process multiple network receive requests simultaneously across multiple CPUs.
- **iSCSI**. iSCSI traffic results in a "double hit" from a CPU overhead perspective. First, the SCSI commands needs to be encapsulated and then this traffic itself needs to be encapsulated so that it can go out on the network. By offloading iSCSI encapsulation, a system can avoid much of this necessary overhead.
- **Distributed Virtual Switch**. The Distributed Virtual Switch is a virtual device that spans multiple vSphere hosts and moves networking management to the cluster level as opposed to managing it at the host level. The Distributed Virtual Switch also enables third party vSwitch integration into vSphere. This component is available only in the Enterprise Plus edition of vSphere 5.1. The switch provides such features as private virtual LANs (PVLANs) and quality of service (QoS). The Distributed Virtual Switch also provides port mirroring capabilities useful in troubleshooting.

Hyper-V

- **SR-IOV**. SR-IOV enables an administrator to directly assign a supported physical NIC to a virtual machine. This option carries with it some network throughout and host CPU performance benefits since the hypervisor does not have to abstract and manage the network communications for that adapter.
- **Large Send Offload (LSO)**. Like vSphere's TCP Segmentation Offload, LSO provides Hyper-V hosts with the ability to submit larger frames – in this case up to 256KB in size – to the network adapter for further processing, thus alleviating the host of that task.
- **Virtual Machine Queue (VMQ)**. Similar to NetQueue. Creates multiple virtual network queues for each virtual machine. Network packets destined for these virtual machines are then sent directly to the VM, reducing some overhead.

- **Hyper-V Extensible Switch**. The Hyper-V Extensible Switch is included in all editions of Hyper-V, including the free edition. This software construct runs in the management partition and brings to Hyper-V a host of capabilities. Most importantly, the software switch greatly simplifies the deployment of Hyper-V into multi-tenant environments by offering private VLANs, DHCP Guard, and improved monitoring capabilities. Microsoft's extensible switch also includes port mirroring capabilities.

Note: Chimney (TCP offload) support, which offloads to the NIC significant portions of the CPU workload normally associated with TCP/IP functionality, has been removed from Hyper-V 2012 due to incompatibilities with Windows Server 2012's NIC teaming capabilities.

Virtual Machine Mobility & Availability

The process of abstracting a running workload from the underlying hardware presents some unique opportunities to shift data and workloads to different hardware platforms. Both VMware and Microsoft provide methods for data and workload migration, but the implementations vary in significant ways.

Workload Migration

A critical availability and maintenance tool, workload migration between hosts, has long been enjoyed by VMware shops (vMotion®) and was introduced to the Microsoft world in previous editions of Hyper-V through a feature the company calls "Live Migration." In previous comparisons, workload migration was one area in which VMware showed its maturity, but Microsoft has effectively closed the gap with Hyper-V 2012, making the two hypervisor products essentially equals in the workload migration category.

VMware vSphere vMotion

vMotion was an incredible introduction to the world of virtualization and revolutionized the way that administrators interacted with active workloads. No longer were services relegated to the original host upon which they were deployed. Now, workloads could be shifted between physical hosts in almost real-time with the end user having no awareness that the activity was taking place.

Although workload transfers via vMotion are initiated with just a few clicks of the mouse, the behind-the-scenes process is quite complex, the technical details being beyond the scope of this paper. Although complex, VMware has continued to innovate in this space and has expanded vMotion's already strong feature set with new capabilities that make this technology even more useful and applicable to more use cases.

In vSphere 5, VMware added two major features to workload migration:

- Use multiple network adapters for vMotion. Allows the VMkernel to transparently load balance vMotion traffic among multiple network interfaces to a point of saturation – even on 10 Gb networks – in order to improve the speed of the transfer.
- Metro vMotion. vMotion has traditionally required that network adapters have a round-trip latency of not more than 5 ms. With vSphere 5, this is increased to 10 ms, making it a suitable solution even in WAN scenarios. This feature is only available in the Enterprise Plus edition.

In vSphere 5.1, which came after the release of Hyper-V 2012, it was VMware playing catch up to Microsoft. VMware added these features:

- Concurrency. Up to eight vMotion operations can take place simultaneously.
- Shared Nothing vMotion/Enhanced vMotion. Prior to the release of vSphere 5.1, in order for vMotion to succeed, there needed to be shared storage between the source and target hosts. With vSphere 5.1, this limitation has been removed, enabling vMotion operations between hosts without shared storage as well as between different vSphere clusters, even if those clusters don't have any shared datastores. However, there are limitations. DRS does not currently support Enhanced vMotion; all vMotion activity that required Enhanced vMotion must be performed manually. Further, only two Enhanced vMotion operations per host are allowed.

Microsoft Live Migration

Workload migration has undergone massive improvement in Hyper-V 2012. In fact, for a period of time, Microsoft leapfrogged VMware in some workload migration technologies, with the latter reacting by catching up with Microsoft in vSphere 5.1. However, Microsoft's workload migration

features still suffer from one negative: The requirement that Hyper-V clusters be built atop Microsoft's clumsy clustering tools.

Storage Migration Comparison

Hardware abstraction applies to both the compute and storage layers of the computing model. This is particularly useful as it provides organizations the freedom to choose among many different storage options and simply migrate workloads between storage manufacturers or performance tiers. Storage migration features add heretofore unheard of levels of simplicity in the otherwise challenging process of upgrading to new storage hardware and keeping virtual machines running at performance tiers appropriate to their function.

VMware Storage vMotion

Storage vMotion provides administrators with a zero downtime storage migration solution that boasts complete transaction integrity. With storage vMotion, virtual machine storage can be migrated in a number of different ways:

- Thick to thin. Migrate virtual disk files from thickly provisioned to thin and back again.
- Raw Device Mapping disk (RDM) to VMDK. Convert an RDM-based disk to a VMDK.
- Across protocols. Migrate NFS-based storage to an iSCSI array or Fibre Channel array and vice versa.

VMware's storage vMotion technique is a no-downtime affair, meaning that it is fully transparent to the end user.

VMware has also added a storage capability to the Distributed Resource Scheduler (DRS), which introduced Storage DRS (SDRS). SDRS automates the leveling and availability of storage resources on an ongoing basis, leaving the administrator to simply create availability and balancing rules and then focusing on more value-add activities.

Microsoft Live Storage Migration

In Hyper-V 2008 R2 SP1 days, Microsoft made a valiant attempt to leverage their Background Intelligent Transfer Service (BITS) technology to migrate live workloads from one storage medium to another, but the service suffered from some major disadvantages when compared to VMware's

Storage vMotion tool. Dubbed Quick Storage Migration, the service was not transparent to the end user, which made the feature particularly poor for production purposes. Those days are gone.

With Hyper-V 2012, Microsoft has completely overhauled their storage migration technology, and as mentioned previously, even briefly leapfrogged VMware's capabilities in this realm. In the latest Hyper-V, Microsoft enables customers to perform unlimited Live Storage Migrations up to the capabilities of the underlying hardware.

Availability Capabilities

Both VMware and Microsoft include high availability capabilities in their products, but they do so to varying degrees. In this area, VMware is currently ahead of the game by providing scalability and features not yet matched in the Hyper-V product. However, Hyper-V provides enough workload availability features as to be relevant in this area.

There are two different scenarios to consider in the availability spectrum. First, there is individual workload availability. What happens when a server or a workload fails? Second, what happens when an entire data center fails? Both are touched upon in this section.

vSphere Availability Mechanisms

On the workload front, vSphere includes the following availability capabilities:

- **VMware High Availability (HA)**. The HA feature monitors virtual machines to detect operating system and hardware failures. When a failure occurs, the virtual machines are restarted on other hosts in the resource pool without the administrator needing to be involved. HA is not fully transparent to the user; the VM needs to be restarted in order for the service to resume operation.
- **Fault Tolerance (FT)**. FT picks up where HA leaves off by providing continuous protection for individual virtual machines that suffer a failure of their host. FT does not protect against application level faults. To work its magic, FT creates a second shadow virtual machine that runs in lockstep with the first. If the original virtual machine's host fails, this shadow virtual machine can assume the workload. The use of FT carries some severe limitations, such as restricting to one the number of vCPUs allowed in the virtual machine.

A plain vSphere implementation doesn't have any host-based site resiliency. However, with the additional of Site Recovery Manager (SRM), an add-on, this capability can be added.

As of vSphere 5.1, a VMware cluster can include up to 32 nodes with 4000 virtual machines.

Hyper-V Availability Mechanisms

In the area of easy high availability, Microsoft remains behind VMware due to Microsoft's reliance on complex clustering when compared against VMFS. However, with Hyper-V 2012 and VMM 2012, administrators can manage Hyper-V clusters completely from within the VMM console, thus providing a bit more simplicity than having to work with multiple tools to manage different parts of the environment. That said, whereas VMware requires the purchase of an additional component to achieve true host- based site failover capability, Hyper-V inherits this capability from Windows Server 2012 in a feature that is known as a stretched cluster. A Windows Server 2012 license also includes MS Site Recovery.

Hyper-V 2012 provides only a high availability capability. At present, an easily implemented solution such as VMware FT is not available, so individual workload protection is not as robust as it can be under vSphere.

Hypervisor Management

Both VMware and Microsoft, in addition to their base hypervisor products, offer robust management toolsets designed to ease the administrative overhead for their products. However, in some instances, the additional management tool layer is required in order to unlock some of the advanced functionality that has been described in this paper. For example, in order to use VMware's Distributed Resource Scheduler to automatically migrate workloads between vSphere hosts, organizations must also acquire the company's vCenter management tool. For Hyper-V, the company's Virtual Machine Manager (VMM) tool is necessary for organizations that want to convert an existing physical server to a virtual machine (P2V).

That is, perhaps, among the biggest differences between VMware and Microsoft in the hypervisor wars. With VMware, it's all but a requirement that organizations that move beyond a single host buy the management product in order to use many features, although both products add a lot to their

respective product's feature sets. With Hyper-V, a lot of functionality – clustering and workload migrations, for example – is provided by the underlying operating system, not the management tool.

This fact is certainly a point in Microsoft's favor. Beyond the basic functionality one would expect from a hypervisor, such as the aforementioned clustering and workload migration features, when one adds VMM, the following processes become possible:

- Fast provisioning of virtual machines
- V2V conversion of VMware-based virtual machines to Hyper-V
- Conversion of physical servers to virtual ones (P2V)
- Template-based virtual machine creation
- Automatic placement of new virtual machines to aid in load balancing
- Centralized management of multiple Hyper-V hosts

Finally, deep performance monitoring, capacity planning, configuration management, and chargeback are beyond the scope of the basic hypervisor management tools and require an additional solution.

Summary

The virtualization possibilities have changed a lot with the release of Hyper-V 2012. No longer is Hyper-V a "good enough" player that provides just a small subset of vSphere's functionality. Indeed, in some ways, Hyper-V has now surpassed vSphere, although vSphere's ease of use still plays heavily in its favor.

About the Author

Scott D. Lowe brings close to twenty years of experience in the IT field, with particular interests in IT management, virtualization, and storage. Lowe is the author or co-author of four books and countless articles and blog posts. Since 2000, he has been a regular contributor to CBS Interactive's TechRepublic unit. Additionally, Lowe is a contributor to virtualizationadmin.com, SearchCIO-Midmarket.com, and Wikibon. When not writing, he produces training videos for TrainSignal.

After spending 10 years as the CIO for a number of organizations, Lowe formed The 1610 Group, an IT strategy and management consultancy focused on assisting clients in discovering value in their IT investments. The 1610 Group operates at both the strategic level with C-level executives and at the tactical level assisting in the implementation of core technologies.

About SolarWinds

SolarWinds (NYSE: SWI) provides powerful and affordable IT management software to customers worldwide from Fortune 500 enterprises to small businesses. In all of our market areas, our approach is consistent. We focus exclusively on IT Pros and strive to eliminate the complexity that they have been forced to accept from traditional enterprise software vendors. SolarWinds delivers on this commitment with unexpected simplicity through products that are easy to find, buy, use and maintain while providing the power to address any IT management problem on any scale. Our solutions are rooted in our deep connection to our user base, which interacts in our online community, **thwack**®, to solve problems, share technology and best practices, and directly participate in our product development process. Learn more today at **solarwinds.com**.

SolarWinds Virtualization Manager

Relying on different tools to manage multiple virtual environments can be a real pain, unless you have SolarWinds in your toolkit. **SolarWinds Virtualization Manager** gives you comprehensive visibility from VM to spindle, through a single pane of glass, into the metrics of your entire virtual infrastructure—including VMware and Hyper-V environments.

Virtualization Manager delivers integrated capacity planning, performance monitoring, VM sprawl control, configuration management, and chargeback automation—all in one incredibly affordable product that's easy to download, deploy, and use.

Download a **free 30-day trial** to try Virtualization Manager in your own environment today.

Learn More

For more information or to purchase SolarWinds products, visit **solarwinds.com** or contact:

Americas
Phone: 866.530.8100
Fax: 512.857.0125
Email: **sales@solarwinds.com**

APAC
Tel: +65 6593 7600
Fax: +65 6593 7601
Email: **sales@solarwinds.com**

EMEA
Phone: +353 21 5002900
Fax: +353 212 380 232
Email: **sales@solarwinds.com**

3711 South MoPac Expressway, Building Two, Austin, Texas 78746