



Why Microsoft?

For Virtualizing & Managing SQL Server

April 2014 | v1.0



Table of Contents

Executive Summary	3
The Cloud OS	4
Scalability & Performance	5
Storage & Networking	9
Flexibility & Resiliency	15
SQL Server Support on Hyper-V	26
Hybrid Infrastructure	27
Comprehensive Management	33
VMware Comparison.....	47
Summary.....	51

Executive Summary

Today's organizations need the ability to seamlessly build, deploy, and manage applications and services across on-premises, private, and public cloud environments. Windows Server, System Center, Microsoft Azure, and SQL Server work together to provide a consistent, integrated platform across the entire IT infrastructure. Throughout this paper, you'll see how the deep relationships between the Windows Server, System Center, Microsoft Azure and SQL Server product groups, ensure that customers benefit from a battle-tested, deeply integrated set of solutions that enable them to get the most from their investments in the platform, and ensure they deliver the best SQL Server service within their IT infrastructures, all at better value, and with more SQL Server-specific benefits, than an equivalent solution from VMware. To summarize:

Scalability & Performance

With support for the largest hypervisor hosts, virtual machines, and failover clusters, Windows Server 2012 R2 provides best in class scalability and proven performance for virtualizing your mission-critical SQL Server workloads.

Storage & Networking

With a comprehensive set of integrated storage and networking capabilities, and support for the latest hardware innovations, Windows Server 2012 R2 provides SQL Server with a flexible, high performing and highly available platform, at considerably reduced costs.

Flexibility & Resiliency

New innovations in migration capabilities provide the SQL Server DBA with complete flexibility for moving mission-critical SQL Server workloads around the virtualized environment, and now, the migrations complete faster than ever. Improvements in Failover Clustering also enhance the availability of SQL Server, whilst providing best in class support for the virtualization of SQL Server AlwaysOn deployments.

Hybrid Infrastructure

Through integration between Windows Server, Microsoft Azure and SQL Server, SQL Server DBA's gain powerful deployment flexibility, enabling rapid deployment of SQL Servers and SQL Server databases directly into the public cloud, or utilizing the public cloud for disaster recovery from on-premises SQL Servers.

Comprehensive Management

With System Center, Microsoft offers best-in-class management for SQL Server. From granular monitoring and powerful time-saving automation, through to data protection and accelerated deployment, System Center enables higher levels of efficiency for managing SQL Server.

The Cloud OS

It's a new day in IT. There are more apps, more devices, and now, more data than ever — all driven by the rise of cloud computing and the use of cloud services. With these technologies playing an ever present role in businesses, how can IT drive more efficiency and deliver new forms of value? Microsoft's answer is the Cloud OS, powered by the technologies you know; Windows Server, System Center and Microsoft Azure.

Big data, the cloud, and bring your own device are converging technology trends that represent real opportunities for IT to deliver more efficiencies and new value. By responding to these opportunities, IT can reduce the cost and complexity of running datacenters at scale, draw insights from any data, support employees wherever they work across any device, and create new business apps or transform existing ones. Enabling all of this is the Cloud OS, our vision of the platform that drives for modern business.

Microsoft's strategy gives IT all the benefits of scale, speed, and agility while still protecting existing investments. This transforms IT into an enabler that can rapidly build and deploy apps, flexibly manage IT services, and support real-time analytics across all forms of data.

This is probably not the first time you've been told about the amazingness of the cloud. You know the pitch: "Do everything you do 100x faster/better/cheaper." And, when used appropriately, the cloud does offer great benefits: increased efficiency that allows IT to focus on strategic business initiatives, a flexible single platform that gives you room to adapt to the future, and centralized IT control even as end-user access expands.

However, all that is called cloud may not be cloud, and not all cloud vendors are alike. But what makes Microsoft different from other vendor offering you the cloud?

Microsoft has both real public cloud scale experience and a deep understanding of your legacy infrastructure. That's the perfect combination to provide the right cloud solution for you. We can adapt to your cloud roadmap and meet your business needs for compliance, security, and high availability. Our hybrid cloud lets you extend your IT team's existing skills, from running on-premises datacenters to Microsoft Azure or a hosting service provider's datacenter. Our experience in datacenters and cloud lets us offer a hybrid cloud that gives you the best of both worlds. Even better, we know that some of the first things you want to move to the cloud may be workloads like SQL Server, SharePoint and Exchange that are critical to your business activity. No other vendor understands Microsoft products and supports them better than we do.

Ultimately, you need a partner who can run your legacy assets, as well as your new cloud assets, on an enterprise level and at enterprise standards. Only Microsoft has the hard-earned experience and expertise across the enterprise cloud landscape that you can bet your business on.

Throughout this paper, you'll learn about many of the capabilities within the Cloud OS, and how, when they come together, they provide the most feature rich, resilient, high performing platform for your key enterprise applications.

To learn more about the Cloud OS, head over to the [Server & Tools website](#).

Scalability & Performance

When it comes to SQL Server, performance is paramount, and with Hyper-V, Microsoft's platform for virtualizing your key workloads, there are a number of capabilities that ensure that you can deliver the highest levels of performance and scale for your mission-critical SQL Server workloads.

Enterprise Class Scalability for running SQL Server

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 data warehousing and analytics 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.




Figure 1: Host, VM & Cluster Scalability in Windows Server 2012 R2 Hyper-V

From an out and out scalability perspective, Hyper-V in Windows Server 2012 R2 greatly expands support for host processors and memory over Windows Server 2008 R2 Hyper-V. When we look at the Hyper-V host itself, as you can see from Figure 1, Hyper-V supports up to **320 Logical Processors** and **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, through aggressive SQL Server consolidation.

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 per VM**, and with the new VHDX virtual hard disk format, support for **64TB per VHDX file**. With each VM supporting over 250 attached virtual disks, it equates to storage in the region of 16PB, per virtual machine. This is huge scale, and opens the door to running high-end, data-intensive, mission-critical, in-memory transactional or analysis workloads that can benefit significantly from that kind of resource capacity. It's important to note, that unlike competitors in this space, this level of scale isn't restricted to a particular edition of Windows Server – on the contrary, it's available in every edition, including our free download, Hyper-V Server 2012 R2. This enables customers to virtualize their largest SQL Server workloads, and experience near native performance when compared with a similarly configured physical SQL Server.

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**. This ensures customers can distribute their mission-critical, virtualized SQL Server workloads across a greater number of physical servers, to provide greater levels of availability yet maintain the single management construct of the cluster. SQL Server also benefits directly from the innovations and improvements in Failover Clustering, as we'll learn later in the paper.


There is another innovation that can drive higher levels of performance for virtualized workloads like SQL Server, and is of particular importance when running virtualized workloads with significant numbers of virtual processors, and high levels of memory. That innovation is Virtual Machine NUMA.

Support for NUMA Architecture

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

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

SQL Server can group its schedulers to map to the grouping of CPUs, based on the hardware NUMA boundary exposed by Windows Server 2012 R2 Hyper-V. For example, a 16-way box may have 4 NUMA nodes, each node having 4 CPUs. This allows for a greater memory locality for that group of schedulers when tasks are processed on the node. From memory perspective, when a thread running on a specific hardware NUMA node allocates memory, the memory manager of SQL Server tries to allocate memory from the memory associated with the NUMA node for locality of reference. Similarly, buffer pool pages are distributed across hardware NUMA nodes. It is more efficient for a thread to access memory from a buffer page that is allocated on the local memory than to access it from foreign memory. And from I/O perspective, each NUMA node has an associated I/O completion port that is used to handle network I/O. This helps optimize the distribution of the network I/O handling across multiple ports.



Guest NUMA topology by default matches host NUMA topology

Figure 2: Physical & Virtual NUMA in Windows Server 2012 R2 Hyper-V


Projecting a **virtual NUMA** topology into a virtual machine provides optimal performance and workload scalability in large virtual machine configurations. It does this by letting the guest operating system and applications such as SQL Server, or the Windows Web Server, IIS, take advantage of their inherent NUMA performance optimizations. The default virtual NUMA topology that is projected into a Hyper-V virtual machine is optimized to match the host's NUMA topology, as shown in figure 2.

Performance of SQL Server on Hyper-V

As we've discussed, Windows Server 2012 R2 Hyper-V can scale to meet the demands of your most intensive workloads. From architectural capabilities such as NUMA, to deep integration with hardware capabilities for powerful offloading, through to Network and Storage QoS. All of these capabilities, and more, help to ensure that when virtualizing your key workloads, they run at their best on Hyper-V.

Over the last 12-24 months, Microsoft has worked closely with Enterprise Strategy Group (ESG), who performed lab testing and analysis on a number of Microsoft key workloads, running virtualized on Windows Server Hyper-V. These included SQL Server, Exchange and SharePoint.


From a SQL Server-specific standpoint, ESG first tested an existing SQL Server 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 graph, over the page.



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

Figure 3 – 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:



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

Figure 4 – Graph of a Hyper-V Enabled SQL Server 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 5 – 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.

Storage & Networking

When it comes to a workload like SQL Server, both the storage and the networking infrastructure are of incredible importance. Windows Server 2012 R2 Hyper-V features a number of advanced storage and networking capabilities and integration points that accelerate performance in these areas, whilst keeping costs low, ensuring it's the optimal platform for running a virtualized SQL Server.

Enhanced Storage Capabilities

Offloaded Data Transfer – Hyper-V Offloaded Data Transfer, or ODX, is a hardware-integrated feature that allows for copying large amounts of data from one location to another. Instead of the Hyper-V host spending valuable CPU and Memory resources to handle a large data copy between 2 locations on the SAN, the task is coordinated and offloaded to the SAN hardware itself. By providing support for offloaded data transfer in the Hyper-V storage stack, this feature makes it possible to complete these operations in a fraction of the time it would have taken without the support. Allowing the virtualized workload, such as SQL Server to use the semantics for offloaded data transfer by passing from the workload to the host hardware helps the virtualized workload operate as efficiently as it would in a non-virtualized environment. The beauty of the ODX solution, is the workload itself need know nothing about ODX itself.

Virtual Fiber Channel - Many enterprises have already invested in Fibre Channel SANs, deploying them within their datacenters to address growing storage requirements. These customers often want the ability to utilize this storage from within their virtual machines instead of having the storage accessible to and used only by the Hyper-

V host. In addition, customers are looking to achieve true SAN line speed from the VMs, to the SAN. Virtual Fibre Channel for Hyper-V provides the guest operating system, and subsequently the workload, such as SQL Server, with unmediated access to a SAN by using a standard World Wide Name (WWN) that is associated with a virtual machine. Hyper-V lets you use Fibre Channel SANs to virtualize workloads that require direct access to SAN logical unit numbers (LUNs). Fibre Channel SANs also let you operate in new scenarios, such as running SQL Server AlwaysOn Failover Cluster Instances inside the guest operating system of a virtual machine connected to shared Fibre Channel storage.

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

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

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

From a SQL Server perspective, more modern versions of SQL Server can also take advantage of these improvements in disk format, as per this [published guidance](#). This means that customers can confidently virtualize workloads like SQL Server, on top of Hyper-V, and utilize the newest underlying disk format for greatest performance, scalability and reliability.

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

- Where the size of the VHD is larger than 2TB.
- To reliably protect against issues for dynamic and differencing disks during power failures.
- To prevent performance degradation issues on the new, large-sector (4KB) physical disks.

As mentioned earlier, Windows Server 2012 Hyper-V introduced a significant update to the VHD format, called VHDX, which has much larger capacity and additional resiliency. VHDX supports up to 64TB of storage. It also provides additional protection against corruption from power failures by logging updates to the VHDX metadata structures, and it prevents performance degradation on large-sector physical disks by optimizing structure alignment. This ensures that SQL Server databases can grow larger than 2TB, fit within a single VHDX file, without needing to manage a more complex disk structure, and as a side benefit, will have increased reliability and performance, especially when running on 4KB physical disks.

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

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

Storage QoS provides the ability to specify a maximum input/output operations per second (IOPS) value for an individual virtual hard disk – this could be, from a SQL Server perspective, the data drive, or the logs/temp drives. Similar to how Network QoS can be used to limit a noisy VM in terms of network traffic and utilization, an administrator can throttle the storage I/O to stop a tenant from consuming excessive storage resources that may impact another tenant. Storage QoS supports Fixed, Dynamic and Differencing virtual hard disks, providing the SQL Server DBA with considerably flexibility.

Lowering Storage Costs for SQL Server

Storage Spaces is a technology in Windows and Windows Server that enables you to virtualize storage by grouping industry-standard, direct-attached disks into storage pools, and then create virtual disks called storage spaces from the available capacity in the storage pools.

Storage Spaces enables cost-effective, highly available, scalable, and flexible storage solutions for business-critical (virtual or physical) deployments. Storage Spaces delivers sophisticated storage virtualization capabilities, which empower customers to use industry-standard storage for single computer and scalable multi-node deployments. It is appropriate for a wide range of customers, including enterprise and cloud hosting companies, which use Windows Server for highly available storage that can cost-effectively grow with demand.




Figure 6 – Windows Server-based Storage Infrastructure

As a starting point, Spaces requires DAS storage, in the form of local or directly attached via SAS, storage. From here, as shown in Figure 6, these physical disks are aggregated or grouped into **Storage Pools**, and from the pool, virtual disks, or **Storage Spaces**, are created. These spaces can have an administrator-defined level of resiliency, from simply **striping** the data across the physical disks, to using **mirroring** or **parity**. These Storage Spaces can

be part of a standalone physical Windows Server file server, or part of a collection of connected Windows Server file servers in a configuration known as a Scale-Out File Server.

Storage Spaces can automatically rebuild mirror and parity spaces in which a disk fails by using dedicated disks that are reserved for replacing failed disks (hot spares), or more rapidly by using spare capacity on other drives in the pool. Storage Spaces also includes background scrubbing and intelligent error correction to allow continuous service availability despite storage component failures. In the event of a power failure or cluster failover, the integrity of data is preserved so that recovery happens quickly and does not result in data loss.

Once a Storage Space is provisioned, it appears to Windows as a disk that can be formatted with an NTFS volume and used for various types of data, such as Virtual Machines, or SQL Server data.

In Windows Server 2012 R2 however, there was additional functionality added to Storage Spaces. As discussed, Storage Pools are the fundamental building blocks for Storage Spaces. With Windows Server 2012, Storage Pools can only consist of a single type of physical disk, i.e. only HDDs, or only SSDs, but not a mix. With Windows Server 2012 R2 however, the administrator can now **combine HDDs and SSDs** in the same Storage Pool, enabling the creation of **Tiered Storage Spaces**.




Figure 7 – Tiered Storage Space in Windows Server 2012 R2

As shown in figure 7, these Tiered Storage Spaces consist of an SSD tier for frequently accessed data, and a HDD tier for less-frequently accessed data. Storage Spaces transparently moves data at a sub-file level between the two tiers based on how frequently data is accessed. As a result, storage tiers can dramatically increase performance for the most used (“hot”) data by moving it to SSD storage, without sacrificing the ability to store large quantities of data on inexpensive HDDs. This is particularly useful for data-intensive SQL Servers which store their data either directly on the Tiered Storage Space, or whether their data resides inside a VM’s VHDX file, on the Tiered Storage Space.

By combining this underlying storage technology, built into Windows Server 2012 R2, across multiple, connected Windows Server 2012 R2 File Servers, IT can create a Clustered File Server, or Scale Out File Server (SOFS), which provides high performance, resilience and a low-cost alternative to traditional SAN storage. To connect to this SOFS, customers would utilize the new SMB 3.0 protocol, which provides high performance, local like speeds, simplified management and control, and reduced costs. As SMB 3.0 is network based, it’s also extremely easy to integrate with your existing network infrastructure for maximum time to value.

Scale-Out-File Server & SMB 3.0 - As stated earlier, to connect to this SOFS, customers would utilize the new SMB 3.0 protocol, which provides high performance, local like speeds, simplified management and control, and

reduced costs. As SMB 3.0 is network based, it's also extremely easy to integrate with your existing network infrastructure for maximum time to value.




Figure 8 – Scale-Out File Server with Hyper-V & SQL Server

However, you wouldn't just want each of your Hyper-V hosts connecting to your SOFS (clustered file servers) via a single network adaptor each. You'd want to provide a level of physical resiliency and increase throughput and performance as much as possible. Windows Server 2012 introduced a new feature called **SMB Multichannel**, part of the SMB 3.0 protocol, which increases the network performance and availability for File Servers. SMB Multichannel allows file servers to use multiple network connections simultaneously and provides the following capabilities:

- **Increased throughput.** The file server can simultaneously transmit more data using multiple connections for high speed network adapters or multiple network adapters.
- **Network Fault Tolerance.** When using multiple network connections at the same time, the clients can continue to work uninterrupted despite the loss of a network connection.
- **Automatic Configuration:** SMB Multichannel automatically discovers the existence of multiple available network paths and dynamically adds connections as required

This would be very beneficial when accessing SQL Server databases from an SMB file share – a new feature that was introduced in SQL Server 2012.

To enhance SMB performance still further, Windows Server 2012 and R2 can harness the power of Remote Device Memory Access (RDMA) capable NICs, through a feature known as **SMB Direct**, to drive even higher levels of performance. SMB Direct, enables NICs to function at full speed with very low latency, while using very little CPU resource. For workloads such as Hyper-V or SQL Server, this enables a remote file server to resemble local storage.

RDMA capable NICs can reach speeds of 56Gbps today, and in the near future, will reach 100Gbps, enabling significant performance advantages over the costly Fibre Channel, or iSCSI alternatives.

Not only can administrators place their virtual machines on these SMB shares, but those running SQL Server have a choice of either placing the important databases and log files inside the VM's virtual disks, or alternatively, providing access from inside the VM, to the SMB shares, and placing the files on the SMB shares themselves, for direct, remote access from SQL Server.

All of these capabilities, combined, provide a high performing, resilient and reliable storage platform for your SQL Server data, placed either directly onto the storage subsystem, or within VMs that reside on the storage. Either way, the solution provide exceptional performance, at considerably lower costs than traditional storage arrays.

For more information on the storage capabilities within Windows Server 2012 R2, you can refer to the [Windows Server 2012 R2 Storage Whitepaper](#).

Enhanced Networking Capabilities

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, such as SQL Server. These capabilities include:

Dynamic Virtual Machine Queue – Dynamic VMQ dynamically distributes incoming virtual machine network traffic processing to host cores (based on processor usage and network load). In times of heavy network load, Dynamic VMQ automatically recruits more cores. In times of light network load, Dynamic VMQ relinquishes those same cores. For network-intensive SQL Server workloads, this increased number of host CPUs helping to process the network traffic, ensures that the network latency remains low, and the SQL Server performance remains high.

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 network devices directly into virtual machines.




Figure 9 – Network Traffic Flows from the VM directly to the SR-IOV capable NIC for increased performance

As shown in Figure 9, this gives VMs the ability to bypass the software-based Hyper-V Virtual Network Switch, and more directly address the underlying physical NIC. As a result, CPU overhead and latency is reduced, with a corresponding rise in throughput, which is very useful for network-intensive virtual workloads such as SQL Server. This is all available, without sacrificing key Hyper-V features such as virtual machine Live Migration.

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. For network-intensive SQL Server instances, vRSS enables increased levels of throughput, and thus, increased performance.

Network QoS – IT administrators also have the ability to control, at a granular level, the network bandwidth allocated to a particular virtual machine. This is important from a workload and application perspective, as you don't want noisy virtual machines to use all of the available bandwidth for a particular application, thus being able to provide a minimal level, and also a maximum level of available bandwidth for a VM, is incredibly useful. It allows IT to provide a consistent level of service for that application or workload, even under contention, and at the same time, restrict certain VMs from consuming all valuable resources from other VMs. From a SQL Server perspective, the SQL Server DBA needs to be able to guarantee a certain level of bandwidth to the SQL Server VM to meet expected service levels. At the same time, if the SQL Server VM has a performance spike, the SQL Server DBA also needs to ensure that the consumption of bandwidth doesn't affect other VMs. This is where Network QoS can help ensure that consistent performance.

For more information on the networking capabilities within Windows Server 2012 R2, you can refer to the [Windows Server 2012 R2 Networking Whitepaper](#).

Flexibility & Resiliency

When it comes to running enterprise workloads, administrators need to ensure that they can take advantage of the underlying infrastructure as best possible, and respond quickly to changes in the environment. In addition, they need to ensure that in the event of a local or geographical disaster, key workloads can be successfully restored in a target location to reduce downtime as much as possible for that workload. Fortunately, Windows Server 2012 R2 provides a number of capabilities in the box that can satisfy those requirements.

Best in Class 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. This includes virtual machines running mission-critical workloads, such as SQL Server. The ability to move virtual machines across Hyper-V hosts was introduced in Windows Server 2008 R2, yet release on release, through to Windows Server 2012 R2, Microsoft have introduced 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. 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, or the contained workloads, such as SQL Server is needed. With Windows Server 2012 R2, customers can migrate VMs faster, taking more advantage of the available network resources, including the ability to simultaneously migrate multiple VMs at once.

SMB-Based Live Migration - As we discussed earlier, 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). From a SQL Server perspective, the DBA gains additional flexibility to move workloads around the Hyper-V environment, without reliance on expensive SAN storage.

Live Migration with Compression - With Windows Server 2012 R2, Live Migration with 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. With this capability, customers can expect to see up to a 50% reduction in time taken to live migrate a running virtual machine between hosts, which is particularly valuable for intensive workloads, like SQL Server, to ensure that the migration completes as quickly as possible.

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. This means customers can take advantage of some of the key SMB features to accelerate the process, such as SMB Multichannel (for resiliency, reliability and increased bandwidth) and SMB Direct, for hardware offloading to RDMA-capable network hardware. This combination provides the most resilient, highest performing live migration capabilities for key virtualized workloads such as SQL Server.

Storage Live Migration - With Storage Live Migration, customers can flexibly move virtual hard disks attached to a running virtual machine, without downtime. Through this feature, you can transfer virtual hard disks, with no downtime, to a new location for upgrading or migrating the underlying storage, performing backend storage maintenance, or redistributing your storage load. From a SQL Server perspective, it may be that the location of the Data VHDX file doesn't currently offer enough IOPS, and a relocation to SSD-based storage will provide increased performance. Storage Live Migration enables the movement of the SQL Server Data VHDX drive to the new location, with no downtime, unlocking the increased performance when moved to the SSD storage.

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

From a SQL Server perspective, this allows the SQL Server DBA to move their SQL Server VMs with complete flexibility, and remove any ties to a particular host, or cluster. For instance, the DBA could be working on a SQL Server VM that is running on a Hyper-V host, with DAS storage, located under the DBA's desk. This SQL Server VM could then be moved onto a production Hyper-V environment, backed by resilient shared storage, live, with

no downtime to the SQL Server VM. These are just some of the flexible migration options that administrators gain through utilizing Shared Nothing Live Migration.

Cross-Version Live Migration - With Windows Server 2012 R2 Hyper-V, customers can upgrade from the previous version, 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.

With all of these migration options, it's never been more flexible for SQL Server administrators to migrate their key workloads between hosts. With Windows Server 2012 R2, Microsoft has made Live Migration operate at incredible speed, ensuring that the time to migrate large workloads is reduced, and the impact on the workload is kept to a minimum.

Integrated Network Virtualization

In order to transform your datacenter, you need to deliver your storage, compute, and networking resources as a shared, elastic resource pool. This datacenter-level abstraction is a critical part of Microsoft's Cloud OS vision.

Based on our experiences running cloud-services like Microsoft Azure, we believe that the network should adapt to application requirements rather than the other way round, which is a limiting factor with traditional networking approaches. Indeed, all infrastructure should be programmable and automated to optimize for application needs, be they around provisioning or operations.

Software Defined Networking, or SDN is all about using software to make your network a pooled, automated resource that can seamlessly extend across cloud boundaries. This allows optimal utilization of your existing physical network infrastructure, agility and flexibility resulting from centralized control, and business-critical workload optimization from deployment of innovative network services.




Figure 10 – A view of the Network Virtualization Layers

To begin, as shown in Figure 10, Windows Server 2012 R2 delivers integrated **Hyper-V Network Virtualization** to help you abstract your applications and workloads from the physical network. With Network Virtualization, IT can dramatically simplify the underlying physical network, with the key goal of ensuring that the physical Hyper-V hosts can communicate on a routable network. On top of this simplified physical network, Windows Server and System Center combine to create VM Networks. These VM Networks provide multi-tenant isolation while running on the shared physical network. This is all done in software, without the need to use VLANs, therefore overcoming the cumbersome management and configuration, along with topology constraints that VLANs typically have. In the example above, you'll note that within the respective blue and red VM networks, the VMs have identical IP

addresses, yet Blue 10.10.10.10 will not conflict with Red 10.10.10.10, thanks to Network Virtualization. To create a new VM network is just a few clicks, meaning creating and managing new environments for DBAs and Developers can be faster, and more responsive than ever, all the while maintaining the desired levels of isolation from other key networking assets. These workloads can also be migrated to any Hyper-V host, and still maintain communication with other VMs that are part of the same VM network, but reside on other Hyper-V hosts. From a SQL Server perspective, imagine a scenario with a service provider, providing SQL Server VMs to customers. With Network Virtualization, SQL Servers can be rapidly deployed for each customer, yet because they will be deployed into isolated VM networks, they won't conflict with one another, and the customers can't connect to other customers' SQL Server VMs – only their own. The same applies in a test/development environment, where SQL Server DBAs and application owners request SQL Servers, but IT don't wish to deploy onto the production network. Network Virtualization provides the isolation from production, but is accessible to the respective DBAs and application owners.

Windows Server and System Center also deliver a **built-in software gateway** that helps bridge physical and virtual networks, thereby enabling frictionless mobility for your applications within and across datacenters. You can also extend your datacenter seamlessly into Microsoft Azure using the Microsoft Azure Virtual Network.

For a more detailed overview of Network Virtualization in Windows Server 2012 R2, please refer to the [Windows Server 2012 R2 Networking Whitepaper](#).

Failover Clustering for VM Resiliency

We've spent a significant amount of time discussing the capabilities of the Windows Server 2012 R2 platform, for providing flexibility for IT admins, DBAs, and consumers of the IT infrastructure. Many of the capabilities that we discussed, focused on scenarios around planned migration or downtime, however what happens in the event of an unplanned outage, such as a loss of an underlying physical host? That is where Failover Clustering, a core component of Windows Server 2012 R2, can provide significant benefits.

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. 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 outage to the VM and its guest OS, 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.




Figure 11 – A Hyper-V Failover Cluster, before and after a loss of a Hyper-V physical node

With Windows Server 2012 and 2012 R2, there have been a number of improvements within Failover Clustering to provide an even more robust and reliable platform for key workloads such as SQL Server. A **reduced dependency, from a cluster perspective, on Active Directory**, helps to provide more deployment and management flexibility, particularly in the event of a failover. Integrated **VM Monitoring** looks inside a particular VM, at key running services, and provides an automated restart capability for that in-VM service, from outside of the VM. From a VM placement perspective, when you have related virtual machines, say, a set of SQL Servers, it may be that you wish for them to run on different underlying physical hosts, to ensure that if a host is lost, it doesn't take down an entire SQL Server VM set – features such as **Anti-Affinity** help with this, and in the event of failover, new **Failover Priority** settings help to ensure that important VMs, such as SQL Server VMs, start before less important VMs. When hosts do fail, thanks to new innovations such as **Dynamic Witness** and **Dynamic Quorum**, Failover Clustering handles node failure in a more optimal manner, helping to ensure that key workloads can be restored on other nodes quickly, even if large numbers of physical nodes go down simultaneously.

All of these capabilities, and many more we haven't discussed, help to ensure that Hyper-V is the optimal platform to run your virtualized SQL Server workloads. Failover Clustering provides a robust and reliable solution that scales as your business grows, and enables you to focus on your workloads, rather than the underlying physical infrastructure.

To learn more about Failover Clustering, and some of the other improvements in this area, within Windows Server 2012 R2, please refer to [TechNet](#).

SQL Server AlwaysOn

Whilst we've spent time discussing the levels of resiliency offered at the physical host level, primarily provided by host-level Failover Clustering, it's important to acknowledge that SQL Server itself, has its own resiliency features and capabilities built in, which provide an application-level of resiliency for your key databases.

SQL Server AlwaysOn was first introduced in SQL Server 2012, with the aim of helping to reduce planned and unplanned downtime through an integrated high availability and disaster recovery solution.

As part of the SQL Server AlwaysOn offering, there are 2 options for ensuring databases are highly available; Availability Groups, which maximizes available for one or more databases and Failover Cluster Instances, which provides redundancy at the server-instance level—a failover cluster instance (FCI).

SQL Server AlwaysOn Failover Cluster Instances - An FCI is a single instance of SQL Server that is installed across Windows Server Failover Clustering (WSFC) nodes and, possibly, across multiple subnets. On the network,

an FCI appears to be an instance of SQL Server running on a single computer, but the FCI provides failover from one WSFC node to another if the current node becomes unavailable.

When there is hardware or software failure of a server, the applications or clients connecting to the server will experience downtime. When a SQL Server instance is configured to be an FCI (instead of a standalone instance), the high availability of that SQL Server instance is protected by the presence of redundant nodes in the FCI. Only one of the nodes in the FCI owns the WSFC resource group at a time.




Figure 12 – SQL Server AlwaysOn Failover Cluster Instance

In case of a failure (hardware failures, operating system failures, application or service failures), or a planned upgrade, the resource group ownership is moved to another WSFC node. This process is transparent to the client or application connecting to SQL Server and this minimize the downtime the application or clients experience during a failure. The following lists some key benefits that SQL Server failover cluster instances provide:

- Protection at the instance level through redundancy
- Automatic failover in the event of a failure (hardware failures, operating system failures, application or service failures)
- Support for a broad array of storage solutions, including WSFC cluster disks (iSCSI, Fiber Channel, and so on) and server message block (SMB) file shares.
- Disaster recovery solution using a multi-subnet FCI or running an FCI-hosted database inside an AlwaysOn availability group. With the new multi-subnet support in Microsoft SQL Server 2012, a multi-subnet FCI no longer requires a virtual LAN, increasing the manageability and security of a multi-subnet FCI.
- Zero reconfiguration of applications and clients during failovers
- Flexible failover policy for granular trigger events for automatic failovers
- Reliable failovers through periodic and detailed health detection using dedicated and persisted connections
- Configurability and predictability in failover time through indirect background checkpoints
- Throttled resource usage during failovers

Contrary to the AlwaysOn Availability Group, an FCI must use shared storage between all nodes of the FCI for database and log storage. The shared storage can be in the form of WSFC cluster disks, disks on a SAN, or file shares on an SMB. This way, all nodes in the FCI have the same view of instance data whenever a failover occurs.

This does mean, however, that the shared storage has the potential of being the single point of failure, and FCI depends on the underlying storage solution to ensure data protection.

In the event of a SQL Server node outage, the running instance is automatically failed over to another SQL Server cluster node. This process doesn't change how the user or application connects to the SQL Server – it's still using the virtual network name for connectivity.

To learn more about SQL Server AlwaysOn Failover Cluster Instances, please refer to [TechNet](#).

SQL Server AlwaysOn – Availability Groups – An alternative to FCI, is an Availability Group. The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, which fail over together. An availability group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.




Figure 13 – SQL Server AlwaysOn Availability Group

As you can see from Figure 13, the Primary Replica of a particular set of databases is on a SQL Server Instance on Node 01 of a Windows Server Failover Cluster. This Primary, is replicated, either in a synchronous-commit mode or an asynchronous-commit mode, to multiple Secondaries. A key difference here, is that the Windows Server Failover Cluster does not require shared storage as it did with the FCI example earlier. Deploying AlwaysOn Availability Groups requires a Windows Server Failover Clustering (WSFC) cluster. Each availability replica of a given availability group must reside on a different node of the same WSFC cluster. The only exception is that while being migrated to another WSFC cluster, an availability group can temporarily straddle two clusters.

For more information on SQL Server AlwaysOn Availability Groups, please refer to [TechNet](#).

As you can see, there are significant advantages to utilizing a SQL Server-level high availability solution. With either SQL Server AlwaysOn Failover Cluster Instances, or Availability Groups, the SQL Server DBA can provide instance or database level resiliency with granular failover control, however, higher levels of availability can be provided when IT combines these 2 levels of availability.

Combining High Availability Solutions

For an extra level of resiliency, customers can combine both Hyper-V host-level failover clustering, in which a node outage results in a VM failover, with a SQL Server-level high availability which provides resiliency in the event of a

SQL Server-level outage. As an example, using SQL Server Failover Cluster Instances as our chosen SQL Server HA solution, to build this combination, we need to construct what's known as a **guest cluster**.




Figure 14 – Combining SQL Server AlwaysOn FCI with Hyper-V Host Clustering before and after Failover

At the Hyper-V level, there are a number of capabilities that enhance this guest cluster approach. Features discussed earlier, such as Failover Priority and Anti-Affinity help restart important SQL Server VMs ahead of others, and also keep the different SQL Server cluster nodes apart on different Hyper-V hosts.

To construct a SQL Server AlwaysOn Failover Cluster inside VMs, you'll need to present some form of shared storage, and from a Windows Server 2012 perspective, that storage could have been iSCSI, direct to the Guest OS over the network, Fibre Channel, using the Virtual Fibre Channel capabilities, or, over the network again, but this time using SMB 3.0 and File Servers as the remote storage. These guest clusters, on Hyper-V, are fully supported with Live Migration, so even though the VMs themselves have an intricate connection, i.e. they are part of a Guest Cluster, IT administrators are fully supported to move those workloads around the Hyper-V cluster to run in the locations that make most sense.

In Windows Server 2012 R2, you could take the same approach, using iSCSI to the guest, Virtual Fibre Channel, or SMB 3.0 to the Guest, however, as an alternative approach, 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 15 – Combining SQL Server AlwaysOn FCI with Hyper-V Host Clustering using Shared VHDX

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.

It's important to note that not only do Microsoft provide a number of options to provide the shared storage for guest clusters, such as SQL Server AlwaysOn Failover Cluster Instances, but they also support Live Migration of those guest cluster nodes, to ensure that the SQL Server DBAs gain the extra levels of resiliency, without sacrificing the agility provided by Live Migration. From a scalability standpoint, guest clusters are restricted to the same node-maximums as in the physical world, enabling the creation of scalable guest clusters on top of the Hyper-V platform.

When we compare this solution with what VMware offer, virtualizing guest clusters is severely crippled in comparison. As documented in the following [VMware KB Article](#), with a shared disk requirement, guest cluster support is limited to a maximum of 5 nodes, and IT admins are required to not use common features such as vMotion (Live Migration) and also DRS (automatic vSphere cluster load balancing using vMotion). Sticking with the shared disk option, whilst VMware do support iSCSI to the guest, Virtual Fibre Channel, and SMB 3.0 through to the guest, if IT want to use a shared VMDK file, to remove the exposure of the underlying physical storage to a VM, the virtual machines that make up the guest cluster have to reside on the same physical vSphere host. If that host suffers an outage, the whole SQL Server Failover Cluster Instance would go down. It's clear from this example that vSphere's solution for guest clustering with shared disks is not ready for production use.

Integrated Cluster Maintenance

There will come a time that the underlying physical infrastructure, that's running your mission critical SQL Server workloads, requires maintenance. This could be in the form of hardware maintenance, or it could be OS-level patches that need to be applied at the Hyper-V layer. The last thing that we want in that scenario, is downtime for those key applications and workloads, but thankfully, with Live Migration, in coordination with Cluster-Aware Updating, we can safely, quickly and easily patch the underlying cluster, whilst keeping the SQL Server workloads running on top.

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.




Figure 16 – Hyper-V Host Level Automated Maintenance with no VM-level Impact

In our example in figure 16, at the host level, CAU will move virtual machines, running SQL Server, around the Hyper-V hosts, using Live Migration, patching the nodes automatically.

Our first Hyper-V host, on the right hand side, is patched, and once complete, one of our SQL Server VMs is automatically moved over to that now-patched Hyper-V node using Live Migration. The now empty, middle node is then patched, and brought up to date. A VM is then migrated to this node from our one remaining unpatched left-hand node. This process, for the hosts, then completes, all with no downtime for the VMs themselves, all the while maintaining our Guest SQL Server Failover Cluster Instance on top.

CAU can also be used to patch the Guest Clusters themselves, providing they are based on Windows Server 2012 or newer. This means that the same interface and tool can coordinate the Hyper-V host cluster, along with the SQL Server Failover Cluster Instance if required.

For further information on patching SQL Server Failover Cluster Instances with CAU, please refer to this [TechNet article](#).

Network Resiliency & VM Replication

To increase reliability and performance in virtualized environments, Windows Server 2012 R2 includes built-in support for **NIC Teaming**. NIC Teaming aggregates physical network adaptors into 'Teams' which provide bandwidth aggregation and failover at a NIC-level. 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.

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

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

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

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

Hyper-V Replica provides asynchronous, software-based 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.

Extended Replication - In Windows Server 2012, Hyper-V Replica would allow replication every 5 minutes, and only between 2 points. So, for instance, a customer could replicate their VMs to a Service Provider, but that would be the furthest that the VM could be replicated. The Service Provider wouldn't easily be able to replicate your VM on to a DR site of their own, for instance.

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

To learn more about Hyper-V Replica, you can visit [TechNet](#).

From a SQL Server perspective, it's important to be aware of the supportability of Hyper-V Replica and the SQL Server workload. SQL Server supports use with Hyper-V Replica, as long as the SQL Server database is not using the SQL Server HA capabilities, such as Mirroring, Availability Groups, Log Shipping, FCI or SQL Replication, and the [published guidelines](#) are followed.

SQL Server Support on Hyper-V

Before customers go ahead and virtualize their SQL Server workloads, it's important to know what is, and isn't supported.

SQL Server Version	Running on the following Windows Server Guest Operating Systems	Supported on Hyper-V	Supports Guest Clustering
SQL Server 2008 SP3	2003 SP2, 2003 R2 SP2, 2008 SP2, 2008 R2 SP1, 2012, 2012 R2	Yes	Yes ¹
SQL Server 2008 R2 SP2	2003 SP2, 2008 SP2, 2008 R2 SP1, 2012, 2012 R2	Yes	Yes ¹
SQL Server 2012 SP1	2008 SP2, 2008 R2 SP1, 2012, 2012 R2	Yes	Yes
SQL Server 2014	2008 SP2, 2008 R2 SP1, 2012, 2012 R2	Yes	Yes

¹Guest Clustering is supported with SQL Server when running a guest operating system of Windows Server 2008 SP2 or newer

Table 1 – SQL Server Support by Version, Operating System and Virtualization/Guest Clustering

As you can see from the top table, a number of SQL Server versions are supported virtualized on Hyper-V, providing you adhere to the supported guest operating system to run that particular version of SQL Server. All 4 of the most recent versions of SQL Server are supported on Hyper-V, and in addition, support being placed into a Guest Cluster on a Hyper-V host(s), with 1 small caveat – the Guest Cluster must be constructed using Guest OS's of Windows Server 2008 SP2 or newer, so a guest cluster running on Windows Server 2003 will not be supported.

SQL Server Version	Supporting Documentation
SQL Server 2008 SP3	Hardware & Software Requirements , Hardware Virtualization & Guest Clustering Support
SQL Server 2008 R2 SP2	Hardware & Software Requirements , Hardware Virtualization & Guest Clustering Support
SQL Server 2012 SP1	Hardware & Software Requirements , Hardware Virtualization & Guest Clustering Support
SQL Server 2014	Hardware & Software Requirements , Hardware Virtualization & Guest Clustering Support

Table 2 – Supporting Documentation for the virtualization of SQL Server by Version

In the bottom table, you'll find links to the respective support statements online, as well as specific virtualization and guest clustering guidance.

Hybrid Infrastructure

Your business requires a flexible IT infrastructure that can scale on demand. With a private cloud in your datacenter, you can be more agile and manage resources more effectively. When you extend your datacenter to meet the public cloud, you are working in a hybrid cloud model. A hybrid cloud gives you the best of both worlds, so you can take advantage of external resources when it makes sense for your business. The Microsoft hybrid cloud combines Microsoft Azure, Windows Server, and Microsoft System Center, giving you Microsoft's enterprise-grade technology in both your company's own datacenter and our global datacenters.

Extending into the cloud should expand your options without adding complexity. A Microsoft hybrid cloud lets you easily move or extend workloads from your datacenter to Microsoft Azure or a hosting service provider's datacenter, while still maintaining a complete view of the infrastructure. You can build hybrid applications that leverage both on-premises and cloud resources. And you can take advantage of storage, backup, and recovery options with increased efficiency and reduced cost.

Integrating with Microsoft Azure

Microsoft Azure is an open and flexible cloud platform that enables you to quickly build, deploy and manage applications across a global network of Microsoft-managed datacenters. You can build applications using any language, tool or framework. And you can integrate your public cloud applications with your existing IT environment. With a 99.95% monthly SLA, Microsoft Azure enables you to build and run highly available applications without focusing on the infrastructure, providing automatic OS and service patching, built in network load balancing and resiliency to hardware failure. Microsoft Azure also supports a deployment model that enables you to upgrade your application without downtime.

Microsoft Azure provides a number of different solutions within its portfolio of services. From web sites, to mobile services, storage to service bus, but one of the more recent services offered within Microsoft Azure is the Infrastructure Services. Microsoft Azure provides on-demand infrastructure that scales and adapts to your changing business needs. Whether you are creating new applications or running existing applications, Microsoft provides best-in-class price-performance and end-to-end support.

These Microsoft Azure-based Virtual Machines deliver on-demand, scalable compute infrastructure when you need to quickly provision resources to meet your growing business needs. With Virtual Machines, you get choice of Windows Server and Linux operating systems in multiple configurations on top of the trustworthy Microsoft Azure foundation. With Microsoft Azure Virtual Machines, there are a number of benefits:

Accelerated Provisioning: Unlock the IT backlog and provision infrastructure at the pace your business requires. Simply choose your compute configuration (standard or high memory instances) and choose an image from the Virtual Machines image gallery. Virtual Machines give your applications and systems full mobility, allowing you to move virtual hard disks (VHDs) back and forth between on-premises and the cloud.

Enterprise Support for Enterprise Products: With 99.95% monthly SLA for multiple instances, Virtual Machines are ready to extend on-premises workloads to the cloud. Microsoft SQL Server, SharePoint Server, BizTalk Server, and many more server applications are validated to run on Virtual Machines. Support for validated workloads is directly provided by Microsoft.

Familiar Tools: If your app runs on Hyper-V, it runs on Microsoft Azure Virtual Machines. This means you can manage your entire infrastructure, on-premises or in the cloud, with System Center. In addition, you can combine your apps running in Virtual Machines with Microsoft Azure platform services - such as Service Bus or Media Services - and extend with new capabilities. With Virtual Machines, you can take advantage of what you already know to achieve new capabilities in the cloud.

Monitoring, Alerting & Scale: Microsoft Azure provides a number of capabilities that help you better understand the health of your applications. You can monitor the health and availability of your applications via health metrics dashboard. You can set up alert rules to be notified when your service availability is degraded. Further, with alerts and notifications, you can define an event of interest that occurs in your application, be notified in real-time when the event occurs and perform actions (manual or automatic) based on those events. Microsoft Azure allows you to configure your application to automatically scale up or down to match the current demands while minimizing costs via auto scale rules.

With Microsoft Azure Virtual Machines, you can deploy your own customized Windows Server or Linux images into a multi-zone, commercially backed SLA production environment in minutes or get started with a pre-configured image from our gallery. With Microsoft Azure and the Virtual Network technology, the cloud is a seamless extension of your datacenter taking full advantage of Microsoft System Center, Active Directory and Visual Studio.

Microsoft Azure Virtual Network enables you to create a logically isolated section in Microsoft Azure and securely connect it to your on-premises datacenter or a single client machine using an IPsec connection. Virtual Network makes it easy for you to take advantage of Microsoft Azure's scalable, on-demand infrastructure while providing connectivity to data and applications on-premises, including systems running on Windows Server, mainframes and UNIX.




Figure 17– A Hybrid Deployment with on-premise & Microsoft Azure-based resources

Now from a SQL Server perspective, DBAs can choose to deploy SQL Server databases directly onto Microsoft Azure SQL Database, without the need for Microsoft Azure based virtual machines. SQL Database is a feature-rich, fully managed relational database service that offers a highly productive experience, incorporates proven SQL Server technology, and delivers business-class capabilities. SQL Database allows customers to scale business apps for burst and global reach by removing the high costs of building an infrastructure which would accommodate occasional peak loads. Customers can also remove the security risks and hassles associated with hosting public-facing apps & websites from within a datacenter by quickly & cost-effectively building websites and mobile &

social apps directly on SQL Database. Finally, customers with strict data policies can get the best of both worlds by keeping sensitive, policy-constrained data on-premises while extending parts of their application to the cloud for added innovation, scale and agility.

Alternatively, for applications that need full SQL Server functionality, Virtual Machines is an ideal solution. You will find SQL Server 2014, SQL Server 2012 and SQL Server 2008 R2 images offered, including standard, web and enterprise editions. If you have an existing SQL Server license with software assurance, as an additional benefit you can move your existing license to Microsoft Azure and simply pay for compute and storage. Running SQL Server in Virtual Machines is a great solution for the following scenarios:

Developing and testing new SQL Server applications quickly – you do not need to wait weeks for on-premises provisioning of hardware, just grab the right SQL Server image from the image gallery and off you go. You can choose to deploy in production or back on-premises with little effort.

Hosting your existing Tier 2 and Tier 3 SQL Server applications – with many VM sizes to choose from and with full SQL Server compatibility you can easily move your existing on-premises SQL Server applications and gain the efficiencies of cloud computing.

Backing up and Restoring your on-premises databases – backup your on-premises database to Microsoft Azure blob storage and have the ability to restore the database to a Microsoft Azure Virtual Machine in case of on-premises disaster recovery.

Extending on-premises Applications – create hybrid applications that utilize on-premises assets as well as Microsoft Azure Virtual Machines for greater efficiency and global reach.

Create Multi-Tiered Cloud Applications – create a multi-tiered application that utilizes the unique scale out capabilities of the SQL Database service for the application tier and leverages the full SQL Server compatibility in Microsoft Azure Virtual Machines for the database tier.

Tuned Images for Data Warehousing – deploy a SQL Server Enterprise image that has been highly optimized for running data warehousing on Azure storage configurations. You can now provision a data warehouse specific image in the cloud within minutes

As stated earlier, DBAs could utilize Microsoft Azure Infrastructure Services by uploading their existing images, uploaded to Microsoft Azure using System Center or using PowerShell. Alternatively, Microsoft Azure also provides a number of pre-constructed templates that customers can utilize to accelerate deployment of OS and application combinations. As you can see from figure 15, there are a number of images to choose from, from Windows Server vanilla images, through to SharePoint, and SQL Server, including SQL Server 2014. This really accelerates the process of deploying SQL Server virtual machines out on to Microsoft Azure.




Figure 18 – Microsoft Azure VM Image Selection Wizard

Another very important attribute of Microsoft Azure Virtual Machines is the fact that they are persistent. That means the state of your Virtual Machine is saved and stored. When you stop your instances and restart later on, your settings and configurations persist and carry forward. When you restart your machine, your settings and configurations persist and carry forward.

In the same context, we need to talk about the geo-replication concept. It is a unique durability benefit Virtual Machines users get on Microsoft Azure. It means you get 6 total copies of your image redundantly replicated.

First, every Microsoft Azure Virtual Machine image is replicated 3 times within the region you are running in. This provides high availability, so if a rack or a storage server were to be down, you still enjoy high availability. You would get a new instance up and running behind the scenes. In addition, since these are persistent Virtual Machines, your new image would have the exact same state and configuration.

Second, your image is asynchronously copied to another datacenter that's hundreds of miles away. These two elements are what we call geo-replication. Geo replication is a unique feature that differentiates Microsoft. Virtual Machine users get that by default. On other cloud platforms, it requires a deliberate effort to capture and image, store it and then to replicate it across regions.

Geo-replication also is key to disaster recovery. Since the images are stored in datacenters hundreds of miles apart redundantly, you can have the peace of mind that you can go to these redundantly stored copies and bring them up in case of natural disasters.

If you are like most other customers, you will be asking if you are locked into the decision to use public cloud and what options you might be giving up when you use Microsoft Azure Virtual Machines. That is pretty straightforward to answer. Since the underlying virtualization format is VHD, and the virtualization layer is Hyper-V, your apps and your images can move in and out of Microsoft Azure and the public cloud. You can start with your app prototypes in the cloud and then move to on-premises when ready or simply move to a hosting partner's datacenter – if so you choose. This is all about flexibility and having options and not being locked-in.




Figure 19 – Rich dashboard interface within Microsoft Azure

Microsoft Azure’s rich web-based portal provides the DBA, or administrator with granular control and visibility into the virtual machines that are running on Microsoft Azure, however for customers who choose to invest in System Center, they can take advantage of utilizing the core set of tools within System Center to also manage their workloads on Microsoft Azure.

Once deployed, SQL Server DBAs can connect to their remote SQL Server instance running in Microsoft Azure, directly from SQL Server Management Studio on-premises, ensuring they can take advantage of both the cloud capacity, scale and redundancy, whilst at the same time, maintain familiarity of tools for working with SQL Server.

SQL Server Hybrid Cloud Availability

Finally, customers can utilize Microsoft Azure to provide a significant increase in the level of availability of SQL Server. If we first think about our on-premises SQL Server deployment, whether on physical machines, or inside virtual machines, SQL Server can provide replication capabilities as part of its built-in high availability capabilities. In more modern versions of SQL Server, this is known as AlwaysOn. AlwaysOn provides Failover Cluster Instances and Availability Groups, which we discussed [earlier](#).




Figure 20 – SQL Server SQL1 replicating to SQL Server SQL2 using synchronous-commit mode, all on-premises

To illustrate further, in the following example, we have a private cloud within the customer’s datacenter. In this datacenter, the customer is operating a number of services, such as Active Directory, but in addition, they have a mission-critical SQL Server workload that is being replicated between two SQL Servers, on-premises. This is using synchronous replication between the primary and the secondary. Before committing transactions, a synchronous-

commit primary replica waits for a synchronous-commit secondary replica to acknowledge that it has finished hardening the log. Synchronous-commit mode ensures that once a given secondary database is synchronized with the primary database, committed transactions are fully protected. This protection comes at the cost of increased transaction latency. This solution gives the DBA resiliency at the database level, but not at the site level.

By integrating with Microsoft Azure Infrastructure Services, the DBA can stand up a SQL Server VM, running in Azure, and run a secondary database in the cloud, backed by the highest levels of availability and scale. Now, before you can add a secondary replica on Microsoft Azure, you'll need to configure a site-to-site Microsoft Azure Virtual Network between your on-premises network and Microsoft Azure, which we discussed earlier.




Figure 21 – SQL Server SQL1 replicating to SQL Server SQL3 using asynchronous-commit mode, from on-premises to Microsoft Azure

Once the connectivity is established and SQL Server configured, SQL Server will replicate up to the Microsoft Azure VM, using the Asynchronous-commit mode. In case of a disaster impacting the on-premises replicas, it's possible to manually failover to the secondary replica on Microsoft Azure, making this the primary replica of the availability group. As part of the failover, the availability group Listener is automatically updated to route connections to the new primary replica on Microsoft Azure.

Besides disaster recovery, the secondary replica(s) on Microsoft Azure can be used to offload reporting applications and backups. This is valuable for companies that require maintaining backups outside of the data center for compliance reasons.

Comprehensive Management

So we've spent a considerable amount of time discussing the virtualizing of the key workloads, along with extending our environment into a hybrid infrastructure, but with System Center, IT admins can benefit from considerable control and efficiency when it comes to deployment, maintenance, and protection.

Now before we start to look at some of the System Center capabilities in more detail, it's important we understand at a high level, the key System Center components that play a part in managing the datacenter. Firstly, we have **Virtual Machine Manager**, which provides the IT administrators the ability to manage the hypervisor hosts, along with the VMs and services on top. It provides complete control of the virtualized infrastructure and centralizes all of the key management tasks and provides significant integration points specifically for SQL Server.

For organizations exploring Private Clouds, self-service becomes an important consideration as a way to provide end users, and application owners with the ability to consume IT infrastructure, in a controlled and abstracted way. **App Controller** provides a common self-service experience that can help you easily configure, deploy, and manage virtual machines and services across private and public clouds. These could be private clouds on your premises, on a service provider premises, or Microsoft Azure public clouds.

We then think about monitoring. How does IT gain insight into the infrastructure, both at the fabric, and the application layers? Well, **Operations Manager** provides infrastructure monitoring that is flexible and cost-effective, helps ensure the predictable performance and availability of vital applications, and offers comprehensive monitoring for your datacenter and cloud, both private and public.

Service Manager provides an integrated platform for automating and adapting your organization's IT service management best practices, such as those found in Microsoft Operations Framework (MOF) and Information Technology Infrastructure Library (ITIL). It provides built-in processes for incident and problem resolution, change control, and asset lifecycle management. It also provides a rich self-service portal for requesting access to resources that could then be consumed through App Controller.

Orchestrator is a workflow management solution for the data center. Orchestrator lets you automate the creation, monitoring, and deployment of resources in your environment. It helps IT to automate key tasks and processes, along with integrating both Microsoft and non-Microsoft systems, platforms and applications, to increase efficiencies within the environment.

Data Protection Manager (DPM) enables disk-based and tape-based data protection and recovery for servers such as SQL Server, Exchange Server, SharePoint, virtual servers, file servers, and support for Windows desktops and laptops. DPM can also centrally manage system state and Bare Metal Recovery (BMR).

Finally, the **Windows Azure Pack** is a collection of Microsoft Azure technologies available to Microsoft customers at no additional cost. Once installed in your datacenter, the Windows Azure Pack integrates with System Center and Windows Server to help provide an alternative self-service portal for managing services such as websites, Virtual Machines, and Service Bus; a portal for administrators to manage resource clouds; scalable web hosting; and more.

Together, these components form System Center. They are available as a single product, licensed via 2 SKUs, Standard or Datacenter, depending on your virtualization needs. For more information on the licensing of System Center, please refer to the [licensing datasheet](#).

SQL Server Monitoring & Insight

Through System Center Operations Manager, DBAs can gain deep visibility into the health, performance and availability of key SQL Server components, such as database engine instances, databases, SQL Server Agents & SQL Clusters, including AlwaysOn.

System Center Operations Manager provides the ability to monitor specific applications and workloads through the importing of a key component known as a management pack. Management packs contain the deep levels of knowledge and monitoring insight on the health, performance and availability of a particular application or workload. This could be a Microsoft-developed workload, such as SQL Server, or a 3rd party workload. With Microsoft workloads, these management packs are developed by Microsoft, however customers and partners can take advantage of a number of tools in order to construct their own management packs to monitor their own specific applications.

The Microsoft-developed management packs are available at no cost, and are designed, engineered and delivered by the teams that craft the workload itself, thus the knowledge and deep level of insight within the System Center Management Pack for SQL Server, has come from the same group designing, developing and delivering SQL Server itself. This is a significant advantage over other monitoring solutions that exist in the industry.




Figure 22 – Key SQL Server components monitored with the System Center Management Pack for SQL Server

The System Center Management Pack (MP) for SQL Server provides System Center Operations Manager with a number of rich capabilities, including the ability to automatically discover the SQL Servers in your environment, proactively monitor the key SQL Server components for health state changes, and provides rich dashboards, infrastructure diagrams and common SQL Server task shortcuts for SQL Server DBAs. As a result, the MP helps the DBAs to quickly identify problems, understand the possible causes, and provide guidance on resolutions. As you can see in figure 22, the MP also includes a number of out of the box views to present data in an easy to

understand and meaningful way to the DBAs within the organization, that will be looking at the information. You will notice that the views are categorized into functional areas to give you a breakdown of the health of a database vs. database mirroring or performance. This allows your DBAs to quickly find the information that they are looking for when troubleshooting a problem.

One of the default views within the MP is the Database State view. This is one that often surprises DBAs as it shows all of the discovered SQL Server databases from all Operations Manager agents deployed in the enterprise. Often, DBAs are unaware of all of the unsupported database instances that exist within their own organization that have been installed for testing or development and often left online after the project was completed and pose a potential security risk within the organization.

Collecting data points to allow you to generate a baseline of normal activity is one of the first steps towards monitoring for performance deviations and exceptions. Without a baseline, you are blind to what to monitor against. The MP for SQL Server includes numerous monitors and rules starting with the basics such as CPU, memory and storage and building from there. The MP for SQL Server has been designed to natively discover and alert on detailed application exceptions and proactively monitor for SQL Server workloads, database statistics and SQL Server locking issues.

The MP for SQL Server also surfaces Policy Based Management (PBM) policies into the Operations Manager console, so users can look at a single pane of glass for availability and performance, as well as configuration issues for custom user policies for AlwaysOn (2012) and Database engine (2008, 2008 R2, 2012). This also offers a centralized view on PBM status across the board.




Figure 23 – SQL Server Summary Dashboard in the System Center Management Pack for SQL Server

As mentioned earlier, the MP for SQL Server provides a number of rich dashboards and views to help articulate key health, performance and availability information for a variety of users. Not only does the MP for SQL Server ship with a number of pre-built dashboards, but the dashboarding capabilities in Operations Manager allow you to quickly build out your own dashboards to display the data that is most relevant to your DBAs. By using Operations Manager default dashboard view, as shown in figure 23, above, users can easily and quickly create a grid view for monitored components. Users can have as many views as they want for their various requirements. The simplicity of creating dashboard views means that your DBA team, with minimal training, can quickly create

and customize their own dashboards to return the information that is most important to them and free the IT administrator up for other tasks.




Figure 24 – SQL Server Service Level Dashboard in the System Center Management Pack for SQL Server

There is a capability within Operations Manager to create SQL Server Service Level Dashboard views that can help you to understand the health, availability and performance of your SQL Server databases and instances. Here you have a quickly-accessible, single pane of glass view into the state of your SLAs, the alerts on your SQL Server instances, the details of your SQL Server instances, metrics on SQL Server transactions per second, recompilations/sec, compilation/sec and batch requests/sec. You are able to easily discover any anomalies and take action on these with the knowledge contained within the MP for SQL Server.

In addition to monitoring, one of System Center Operations Manager's key strengths is reporting, which provides a historical look at health, performance and availability, using data stored in Operations Manager's vast SQL Server-powered Data Warehouse. The out of the box reporting included within the MP for SQL Server is also extensive. If there isn't a specific report included in the MP for SQL Server, you can open one of the reports in the Generic Report Library and add the SQL Server classes to the report that you wish to report on. There are a number of reports included within the MP for SQL Server, ranging from database space to lock analysis, and from user activity to the top 5 deadlocked databases.

For more detail on the System Center Management Pack for SQL Server, please visit the [download site](#).

SQL Server Configuration Guidance

With System Center Advisor, both DBAs and IT administrators alike can benefit from a free cloud based service, which provides best practice configuration guidance for both SQL Server, Windows Server, Hyper-V and a number of other workloads.

System Center Advisor will proactively scan your chosen servers, and compare the current configuration with its huge database of best practice guidance and configuration optimizations, built up from literally thousands of hours of customers support engagements. From this, it will suggest specific configuration changes to optimize performance. Changes could be related to a missing patch, a specific application flag, or an OS setting, to name

but a few. These configuration alerts are centralized either into the web console, or, for customers with an investment in Operations Manager, centralized right within Operations Manager, alongside health and performance alerts for your other applications and services.




Figure 25 – System Center Advisor providing SQL Server specific configuration alerts

When a particular alert is selected, you benefit from detailed, focused knowledge and remediation guidance that is specific to the workload in question, in this case, SQL Server. In addition, where appropriate, you're provided links to relevant KB articles that may have been published that relate to this specific problem. As Advisor is using a live database of resources and guidance, which is being added to all the time, you'll find more relevant information than you'll find just through KBs online.




Figure 26 – System Center Advisor providing SQL Server specific configuration alert

All of this helps to ensure that your infrastructure is optimally configured, any problems that do arise, can be resolved quickly, and efficiently. For more information on System Center Advisor, please visit the [System Center Advisor landing page](#).

SQL Server Automation

With System Center Orchestrator, IT admins, along with the DBA's themselves, can save time and effort, along with providing improved standardization and consistency, through the use of automated runbooks. For the SQL Server DBA in particular, there are a number of manual tasks that eat into valuable time, from a patching process, through to a more general set of SQL Server actions that perhaps require a certain order to execute.




Figure 27 – An example runbook for automating SQL Server database indexing with Orchestrator

A common task, as shown in figure 27, is the management of indexes. Here, the DBA, or IT admin, constructs a simple runbook which first asks for some data; perhaps a certain user who owns a set of databases. Based on this user, Orchestrator talks to SQL Server to retrieve a list of databases, and for each DB, triggers an activity that will evaluate the indexes of each database. Should it determine that the indexes are fine, Orchestrator will execute along the bottom path, issuing an alert, or email message to the specific user to let them know all is fine. Alternatively, should indexes need rebuilding or reorganizing, that particular branch of the runbook is executed, automatically, and again, the user is notified about the end result.

This is just one example of a powerful, yet simple runbook, which automates a usually-time consuming and manual task, and highlights the value of Orchestrator not only for SQL Server, but as a more broadly utilized solution to lower the total cost of ownership for running the IT environment. For more information on System Center Orchestrator, visit [TechNet](#).

SQL Server Protection

We've spent a considerable amount of time looking at how we can monitor, configure and automate SQL Server, however, what happens in the event of a failure? We talked about clustering earlier, and also Hyper-V Replica as solution to provide a more real-time solution for hardware or OS failure, but what about if we need to restore SQL Server after an unrecoverable loss? For that, we need to think about backup.

With Data Protection Manager, Microsoft provides a solution to protect key SQL Server workloads, in addition to SharePoint, Exchange, and Hyper-V VMs, along with the core Windows Server OS, all from the same single agent. When running SQL Server inside a Hyper-V virtual machine, not only can the administrator protect at the VM level, working with both the Hyper-V and SQL Server VSS writers to coordinate the backup, but additionally, IT administrators can place an agent inside the SQL Server VMs and gain a more granular, centralized level of backup for the SQL Servers too.




Figure 28 – DPM Protecting SQL Server Instances & Datacenter

This backup can be made to short term disk, long term tape, or even, through integration with Microsoft Azure, to the cloud for off-site storage and protection. Non-Microsoft backup solutions tend to take generic backup functionality and adapt it to support-specific applications. In contrast, Microsoft created DPM to leverage fully supported Microsoft technologies in order to provide near-continuous data protection including Windows Server hosts running SQL Server. This support includes:

- SQL Server Instance-level protection
- Data source co-location, which allows a DPM server to protect over 2,000 databases
- Database Administrator Self-Service Recovery
- Backup for SQL Server high availability database configurations
- Data recovery at database level

DPM is engineered to provide the best possible protection for SQL Server. In fact, the team that built DPM worked in consultation with the team that built SQL Server to ensure that SQL Server workloads are reliably protected. DPM seamlessly interacts in the following ways with the SQL Server VSS writer to capture consistent versions of a SQL Server deployment without interrupting access databases:

- A baseline copy of the SQL Server data can be made using either the DPM block-level synchronization engine or can be done manually.
- Express full backups are captured. These backups use the SQL Server VSS writer and DPM agent to identify which blocks have changed in the database and then forward those changed blocks to the DPM server.
- Database transaction logs are synchronized with the DPM server as frequently as every 15 minutes between express full backups. DPM synchronizes the log files using a VSS incremental operation.
- Best practice is to configure express full backups every evening and more often for transaction log synchronization (between every 15 to 60 minutes).

DPM provides the administrator with flexibility around restoration, supporting restoration to the original, or a different SQL Server, and in addition, the Self-Service Recovery Tool (SSRT) for System Center 2012

R2 – Data Protection Manager (DPM) enables end users to recover SQL Server databases that are backed up by the DPM server, without any intervention from the DPM administrator.

Once a SQL Server Instance has been protected by DPM, any databases that are added to that instance can be automatically protected, ensuring that the administrator can provide a reliable service to the DBAs yet in an automated fashion. For more information and guidance on System Center Data Protection Manager, please refer to the [DPM documentation](#).

Delivering SQL Server as a Service with Virtual Machine Manager, Service Manager & Orchestrator

Aside from monitoring, automating and protecting, System Center can also provide considerable value when it comes to provisioning key workloads like SQL Server, in particular, delivering SQL Server as a Service, i.e. SQL Server VMs or Databases, or even instances, on demand, driven by self-service. To provide this functionality, we can combine a number of System Center components; Virtual Machine Manager, Service Manager & Orchestrator.

Firstly, with Service Manager, we're able to construct Service Offerings that reflect a menu of choices that IT would like to make available to DBAs to consume. As we can see from the figure 29, we have a choice of 3 Service Offerings, and these are accessed by the DBA via the Service Manager self-service portal. These offerings that we have provided to the DBAs are as follows:



Figure 29 – Service Offering in Service Manager to request SQL Server Databases/VMs

- **New SQL Server Database** - Request a new SQL Server DB on a shared or dedicated Instance
- **New SQL Server Virtual Environment** - Request a complete new set of SQL Server-enabled virtual machines, which are ready for the DBA to utilize exclusively
- **New SQL Server Instance** - Request a new SQL Server instance on shared or dedicated server, on which Databases can then be deployed.

Request New Database

Please fill the requested information. All fields are mandatory.

Database Name
DB001

Database Size
100

Shared or Dedicated
Dedicated

High Availability?

← Back Next → Cancel ✕

Figure 30 – Request Offering within Service Manager to request a new Database

If the DBA were to select the 'Request New Database' Service Offering, the DBA is able to provide some very basic details, such as database name, database size, and whether it needs to be deployed in a shared environment, or dedicated. Each of these choices will reflect a different end result, which we'll learn more about later.

The DBA also has the ability to select whether the Databases, or Instances, can be deployed on a standalone (not highly available) SQL Server, or whether the database/instance will be deployed into a clustered SQL Server through the 'High Availability' tick box. These Service Offerings are just examples of what can be provided. These can be customized to reflect your specific environment.

In addition, Service Manager contains approval mechanisms that allow the DBA to request resources, and only on approval, will the VM, Database or Instance be deployed into the environment using automation. This way, IT has complete control of delivery of the SQL infrastructure, and in addition, through automation.




Figure 31 – End result of a submitted request – Details of the database that has been created for the DBA

Finally, the DBA will receive an email upon completion of the request, containing all the relevant information that the DBA requires to connect and start to administer that particular SQL Server, Instance or Database.

Once the DBA has completed their request through the Service Manager self-service portal, System Center Orchestrator will take the relevant information that the user has submitted, and begin the creation/deployment of the new SQL Server Database, VM or Instance.




Figure 32 – Example Runbooks that automate the activities presented via the Service Manager Service & Request Offerings

Orchestrator handles this through runbooks – standardized processes that step through specific tasks. Not only will Orchestrator automate key processes, such as the deployment of a VM, or a new Database, but to do so, it connects systems, both Microsoft and 3rd party. You’ll see from figure 32, there are a number of branches that the runbook can take, all based on the original form that the DBA completed, and each of the branches themselves, call other runbooks to perform the specific tasks, such as creating a new SQL Server instance, or database on an existing instance.

This runbook, and it’s connected runbooks, whilst appearing complex on the surface, simply break down the overall tasks into a series of steps, all of which make decisions based on the original user input, but what this provides to IT, is the ability to standardize the approach to rolling out SQL Servers into the environment to ensure a consistent, well managed SQL Server infrastructure.

Whilst the Orchestrator runbook will communicate directly with the virtualized SQL Server to deploy a new instance or database onto an existing SQL Server, there are a number of scenarios that will require the deployment of a new SQL Server VM in its entirety, and to accelerate this, Virtual Machine Manager allows the construction of granular templates, that allow a standardized deployment of virtual machines into the environment.




Figure 33 – Single Tier Service Template for deploying SQL Server VMs in Virtual Machine Manager

In this example, in figure 33, we have a single tier Service Template, which defines the virtual hardware, such as vCPUs, memory, disk etc., along with specific guest OS and application-level configuration items. Whilst this is a

single tier Service Template, VMM supports multi-tiered Service Templates, for deploying more complex applications such as SharePoint.

As stated, the template contains a hardware profile, which defines the virtual hardware, and every VM based on this template will look, at least from the hardware characteristics perspective, the same. The Template also contains an application configuration profile, which allows the administrator to add on specific app-level customizations. These could be in the form of scripts, or workload specific configuration items such as SQL Server DAC packages, or Web Deploy packages.




Figure 34 – Adding SQL Server specific configuration information for a VM template in Virtual Machine Manager

As we can see here, this particular example shows a specific SQL Server configuration, which allows the IT administrator to include information from SQL Server administrators, to product key information, and much more.

Finally, when Orchestrator reaches the correct point in its runbook, it will communicate with VMM, coordinating VMM to deploy, in this particular example, a new SQL Server VM to a set of existing, dedicated instances. VMM will deploy the VM from its template library, so each new one is the same as the last. The process is fully automated, and consistent.

Name	Status
SQL Farm 01	Deploying
SQL Server Instances VMs	OK
DEMO-WS-SQL-1.ad.corp.local	Running
DEMO-WS-SQL-2.ad.corp.local	Running
DEMO-WS-SQL-3.ad.corp.local	Creating...

Figure 35 – A new SQL Server VM instance being deployed via Virtual Machine Manager

Once deployed, the SQL Server DBAs have a number of options in order to communicate, and integrate with that SQL Server. IT may choose to provide direct RDP access to the virtualized OS, or alternatively, they may simply use their local SQL Server Management Studio to connect to the remote SQL Server instance over the network. There may be times however, where IT wants to isolate these virtualized SQL Server workloads, upon which the use of a secure, rich, self-service portal, can be very useful.

With System Center App Controller, IT can provide delegated, controlled access to a rich, Silverlight based portal, where the DBAs can only see, and access the resources they have been delegated.

In this case, as you can see from the slide, the DBA only has access to their specific SQL Server cloud, of which they have a finite amount of capacity to utilize for their SQL Server VMs. This allows IT to efficiently control usage of the underlying fabric and presents the DBA with a clear representation of what they have, and what they are using.




Figure 36 – A Cloud as seen through App Controller

Within the clouds, the DBA can deploy as many VMs, typically from templates, as they can fit within their capacity limit, but IT can also restrict the actions that the DBAs can perform against those VMs. For instance, a DBA may be allowed to deploy new VMs, from templates, but they may not be able to stop or delete VMs.

Finally, the DBA has the ability to console access their VM, without the VM needing to be on the corporate management network, which means the IT team can maintain their isolation, and the DBA can have a rich, full fidelity experience when administering their SQL Server.

To download the 'SQL Server as a Service Kit', which uses System Center Virtual Machine Manager, Service Manager and Orchestrator, [visit the Building Clouds blog](#).

Delivering SQL Server as a Service with the Windows Azure Pack

An alternative approach to using the previously discussed System Center Service Manager, Orchestrator and Virtual Machine Manager combination, with the 'SQL Server as a Service' kit, would be to harness the power and built-in capabilities within the Windows Azure Pack.

As mentioned earlier, the Windows Azure Pack is a collection of Microsoft Azure technologies available to Microsoft customers at no additional cost. Once installed in your datacenter, the Windows Azure Pack integrates with System Center and Windows Server to help provide an alternative, rich and intuitive self-service portal for

managing services such as websites, Virtual Machines, Databases and Service Bus; a portal for administrators to manage resource clouds; scalable web hosting; and more.

When it comes to SQL Server specifically, the Windows Azure Pack can provide a number of attractive options for the provisioning and consumption of SQL Servers.

SQL Server via Infrastructure as a Service (IaaS)

One of the fundamental capabilities provided by the Windows Azure Pack, is Infrastructure as a Service (IaaS). As an IT Administrator, you can use System Center 2012 R2 Virtual Machine Manager in combination with the Windows Azure Pack to provide your tenants with the ability create virtual machines with specific operating systems and applications already installed, and thus deliver to them, Infrastructure as a Service.

To enable this, the IT Administrator can create Virtual Machine Role templates. Tenants, or end users, can use these templates to create Virtual Machine Roles for deployment into the infrastructure. To be more specific, think about an IT Administrator creating a SQL Server VM Role Template. This VM Role Template would not only contain the appropriate guest operating system, such as Windows Server 2012 R2, but in addition, would have a specific configuration attached to it, to enable certain roles and features of the operating system at deployment time, and would also contain the appropriate files and configuration data to perform a fresh deployment of SQL Server, on top of this guest operating system. Each time a SQL Server DBA requests a new 'SQL Server VM Role' from this template, the Windows Azure Pack, in conjunction with System Center, orchestrates the deployment in a completely automated fashion, based on minimal data entry from the SQL Server DBA. The end result? The SQL Server DBA is provisioned with a new SQL Server-enabled virtual machine, in which they can access, and start working with, without having to install and configure SQL Server themselves.




Figure 37 – Gallery Items in the Windows Azure Pack

These templates are presented to the end user, in this case, the DBA, within the Windows Azure Pack, in the form of a gallery. This gallery lists the templates that have been made available to that particular user, and as you can see in figure 36, this particular user not only has the ability to deploy standard Windows Server images, but also customized (IIS, AD) Windows Servers, along with Windows Servers with particular applications configured within the image, such as Exchange, SharePoint, Lync, SQL Server, and even Oracle.

The DBA can complete a simple wizard, entering information relating to the desired virtual machine itself, such as name, network, domain join etc., along with specific workload-related information, such as SQL Server instance name, initial database creation, SQL Server administrators, and more. Upon clicking submit, the Windows Azure Pack, along with System Center 2012 R2 Virtual Machine Manager, coordinate, orchestrate and automate the quick and streamlined deployment of that new SQL Server image for the DBA. To accelerate population of your VM Role gallery, Microsoft publishes a portfolio of Microsoft and Partner-produced gallery items that can be centrally downloaded for use in the Windows Azure Pack via the Microsoft Web Platform Installer.




Figure 38 – Monitoring the performance of the deployed workload within the Windows Azure Pack

Once the SQL Server VM Role has been deployed, the SQL Server DBA has the ability to view performance metrics and usage data for their respective VMs, along with a number of VM-related controls, such as starting and stopping, restarting, and also, connecting to the VM, either via RDP over the network, or if the VM is operating in a disconnected or isolated environment, using the Hyper-V console connection capability to administer the OS, and SQL Server instance.

Database as a Service with SQL Server

Earlier, we discussed Microsoft Azure SQL Database, which provides DBAs the ability to deploy SQL Server databases directly onto the SQL Database platform, without the need for Microsoft Azure based virtual machines (IaaS). That same functionality is available for consumption on-premises, within the Windows Azure Pack.




Figure 39 – SQL Server Databases that have been provisioned within the Windows Azure Pack

The IT Administrator can construct a SQL Server-based platform at the backend, enabling key SQL Server capabilities such as SQL Server AlwaysOn, for resilience and availability, and expose the SQL Server platform to the DBAs and application owners, that's simple and quick to consume and provides a rapid deployment mechanism to meet their needs. The DBAs and application owners who don't need access to the SQL Server instance configuration, or the underlying operating system, would be well suited to this consumption model.




Figure 40 – Key SQL Server database connectivity information presented via the Windows Azure Pack

With a new database deployed, via self-service, the user can quickly connect to this database, using the information provided by the Windows Azure Pack, directly from their common management tools, such as SQL Server Management Studio.

Both of these options, provided by the Windows Azure Pack, enable IT to deliver a standardized set of offerings to their users and consumers, with enough flexibility to meet the needs of specific groups. For users who need greater granular control, the IaaS capabilities, with SQL Server-enabled VM roles, is a preferred option, whereas for those users who simply need a 'database', the DaaS offering with the Windows Azure Pack provides them with an excellent solution.

For more information on the Windows Azure Pack, read the [Windows Azure Pack whitepaper](#).

VMware Comparison

Throughout this paper, we've seen a number of compelling reasons as to why Windows Server 2012 R2, System Center 2012 R2 and Microsoft Azure, provide the optimal platform for running key, mission-critical workloads like SQL Server. From the scalability, performance and rich enterprise feature-set offered by Hyper-V, to the comprehensive management capabilities for deploying, managing and maintaining SQL Server with System Center, and on to seamless integration with Microsoft Azure, it's clear that the Microsoft platform offers considerably more than key competitors. When we compare specifically with VMware, this becomes even more evident.

Virtualization Comparison

Virtualization Platform Capabilities	Microsoft	vmware
Enterprise-class scalability in all editions	✓	✗
Proven performance for virtualizing latest SQL Server editions	✓	✗
Production-ready storage capabilities to lower costs	✓	✗
Network performance enhancements without restrictions	✓	✗
Flexible VM migration with accelerated performance	✓	✗
Integrated Network Virtualization without additional costs	✓	✗
Comprehensive host & guest clustering support	✓	~
VM replication with near-synchronous replication	✓	✗

Table 3 – Comparing Hyper-V and VMware vSphere for Virtualizing SQL Server

Only Hyper-V provides the enterprise class scalability, such as support for 64 vCPUs, in all editions, unlike VMware, which restrict that luxury to only the highest edition, vSphere 5.5 Enterprise Plus. The scalability Hyper-V provides ensures that customers can run their largest scale up, and scale out workloads with confidence.

We at Microsoft also know it's important to prove that our enterprise workloads run well, in fact, run best on the Microsoft platform, which is why we work with 3rd parties to rigorously test, and [publish performance reports](#) that give our customers confidence in virtualizing the workloads. We don't rely on aging reports like VMware, who have reports running legacy versions of the workload, on legacy versions of their platform, whereas with Microsoft, we provide the latest shipping workloads, on the latest platforms, to ensure you can make a decision based on the most up to date and relevant information.

With Windows Server 2012 R2, Microsoft is also bringing to market solutions that allow customers to re-evaluate storage, particularly for enterprise workloads. With Storage Spaces, Tiering, SMB 3.0, SMB Multichannel, SMB Direct, and the new Scale-Out-File-Server, Microsoft is bringing to market a solution that provides incredible performance, combined with the simplicity and flexibility of management. VMware's equivalent storage virtualization capabilities are unproven, and primarily aimed at Tier-2/Tier-3 workloads.

From a networking perspective, again, Microsoft's key network performance capabilities are in the box, without the need for the highest SKU, expensive add-ons or additional products. Our performance-enhancing network capabilities, like SR-IOV, can provide a significant boost to performance of workloads like SQL Server, yet don't mean you have to lose agility in the form of Live Migration. That's something that can't be said of VMware, who, when SR-IOV is enabled, force customers to sacrifice vMotion.

At Microsoft, we're always striving to make things faster, more efficient, and deliver higher performance, and whilst VMware stopped innovating around their vMotion technology, we knew we could provide more capability, better performance for Live Migration to reduce the impact on the running workload. Live Migration with Compression, and Live Migration over SMB, enhanced with SMB Direct are perfect examples of those enhancements to which VMware has no equivalent, and best of all, it's included in all editions of Hyper-V, meaning organizations large and small can take advantage.

We also touched on Network Virtualization, as a mechanism to take control of the network, manage through software, and deliver new networks, for isolation, development & test etc., on demand. With Hyper-V Network Virtualization, everything you need is in the box with Windows Server & System Center. With VMware, you need

to navigate their portfolio like a minefield, and then decide how to connect the dots of acquired and in-house developed technologies. With Microsoft, our strategy is clear, and benefits can be realized quickly.

We spent considerable time discussing how to provide resilience for SQL Server. Through Failover Clustering, Microsoft provides resilience for a virtualized SQL Server infrastructure at 2 levels – firstly with the host, being able to failover SQL Server virtual machines in the event of hardware failure, but secondly, through the integration between SQL Server and Windows Server clustering, SQL Server AlwaysOn Failover Cluster Instances takes advantage of all of Windows Server’s new Failover Clustering capabilities, such as support for 64 cluster nodes. Microsoft provides the most flexibility and support of Guest Clusters i.e. SQL Server AlwaysOn Failover Cluster instances, on top of a Hyper-V Cluster. Whilst VMware do support guest clustering, they lack support for more than 5 SQL Server nodes in a virtualized cluster, and restrict the agility of customers by not supporting vMotion of guest cluster nodes. No such restrictions exist with the Microsoft platform. In addition, when using Shared Virtual Disks, VMware restrict the Guest Cluster to operating all on the same host, defeating the object of high availability.

Finally, only Microsoft’s platform, with Windows Server 2012 R2 Hyper-V, provides in box VM replication to a near synchronous level, with replication times down to every 30 seconds. With VMware, their in-box, software based replication is only as quick as every 15 minutes, which in SQL Server terms, could be a lifetime.

For a more detailed, feature level comparison of Hyper-V & vSphere, please see the [Competitive Advantages of Windows Server 2012 R2 Hyper-V over VMware vSphere 5.5 whitepaper](#).

Hybrid Comparison

Hybrid Capabilities

Cloud-based SQL Server relational database service
 Public cloud integration within SQL Server
 Public cloud image gallery with SQL Server images

 Microsoft	vmware
	
	
	

Table 4 – Comparing Hybrid/Public Cloud Integration with SQL Server

With Microsoft Azure, Microsoft provides a relational database service called SQL Database, enabling you to easily provision and deploy relational database solutions. Benefits include manageability, high availability, scalability, a familiar development model, and a relational data model. VMware however, did have a largely-unknown database technology, vFabric SQLFire, until the split of technologies, resulting in the Pivotal company taking ownership of these, yet it’s not offered within their vCloud Hybrid Services as a platform to run your databases so customers are restricted solely to running inside virtual machines, which doesn’t always fit with every scenario.

Additionally, as Windows Server, SQL Server, and Microsoft Azure are engineered by Microsoft, it means specific integration points can be built right into the products. An example of this is the integration within SQL Server 2012 for backing up SQL Server databases, to Microsoft Azure, baked right into the SQL Server Management Studio UI. VMware on the other hand, have no specific functionality exposed around SQL Server and public cloud in either vSphere, or their vCloud Hybrid Services. In addition, as we discussed earlier, SQL Server also provides replication capabilities to asynchronously replicate data to SQL Servers running in Azure.

Finally, Both VMware and Microsoft provide a cloud service with a number of pre-built images for accelerated deployment. On the VMware side, they do provide a few SQL Server-specific images, however not as many as

Microsoft, which for instance, already offers customers the ability to test SQL Server 2014, in Microsoft Azure, allowing customers to quickly test and evaluate, without needing to install and configure themselves.

Management Comparison

Management Capabilities	Microsoft	vmware
SQL Server monitoring built by SQL Server Product Group	✓	✗
Cloud-based SQL Server configuration guidance	✓	✗
Orchestration & automation of SQL Server	✓	~
Granular SQL Server protection to disk, tape & cloud	✓	✗
Integrated Service Management for SQL Server as a Service	✓	~
SQL Server specific virtual machine templates	✓	~
Rich, intuitive self-service consumption of SQL Server VMs	✓	~

Table 5 – Comparing Management Capabilities for SQL Server

Firstly, we discussed monitoring, and the powerful role that Operations Manager can provide to gather deep and meaningful information about the entire infrastructure, along with SQL Server specific metrics and health information. Only Microsoft provides a free management pack, to gather this SQL Server-specific information, and these management packs are built by the teams that engineer the SQL Server product. VMware however, rely on their own experience of SQL Server in order to provide their basic monitoring capabilities and metrics. This doesn't compare well with Microsoft's deep SQL Server insight, along with rich presentation of SQL Server-focused information through views and dashboards.

With System Center Advisor, Microsoft provides a cloud service that helps you to run your key workloads more efficiently, optimizing their configuration to avoid potential problems in the future. Granted, as this is at the OS level, this would work for a SQL Server VM running on vSphere also, however, it would provide no benefit to the vSphere hosts, but would help to avoid mis-configuration within the underlying Hyper-V environment. VMware offer now similar level of service for your key workloads such as SQL Server.

With System Center Orchestrator, Microsoft provides the ability to construct powerful, SQL Server-specific runbooks to automate repetitive manual actions and tasks, freeing up valuable time for the DBA. Microsoft, partners such as Keverion, and the community have constructed powerful Integration Packs that extend the inbox Orchestrator SQL Server-capabilities even further, to enable customers to construct even more comprehensive runbooks and workflows to automate SQL Server. VMware on the other hand, rely on very limited set of pre-defined SQL Server automation capabilities for their vCenter Orchestrator component, which are significantly less than that provided inbox by Microsoft. In addition, this SQL Server Plug-in for vCenter Orchestrator has not seen a significant update since June 2012.

With Data Protection Manager, customers have the ability to protect their SQL Server VMs, and SQL Server databases, at a granular level, to disk, tape and cloud. Whilst VMware, with its new Data Protection Advanced SKU, which, itself, is an additional costly purchase on top of vSphere, they do provide SQL Server-level protection, it's only to disk, with no integration with tape, or public cloud services like Microsoft Azure.

When it comes to deploying SQL Server, we looked at how the combination of Service Manager, Orchestrator and Virtual Machine Manager combine with the [SQL Server as a Service kit](#), that Microsoft provide to customers at no charge, to enable a request-based, ITIL aligned, automated SQL Server deployment solution that not only allows the DBA to deploy full virtual machines with SQL Server enabled, but also, deploy new databases and instances

onto existing SQL Servers and SQL Server Clusters. Alternatively, customers can also take advantage of the Windows Azure Pack, as an alternative way to provision SQL Server-based virtual machines, or alternatively, SQL Server Databases through the powerful DaaS capabilities.

Whilst VMware have some of the tools to enable this kind of functionality, it's all left to the customer to construct. With vCenter, the template granularity has no workload or application specific focus, and with their vCenter Orchestrator, VMware rely on very limited set of pre-defined SQL Server automation capabilities, which are significantly less than that provided in-box by Microsoft. From a Service Management perspective, the core vSphere product line doesn't provide the IT admin with the ability to allow a DBA to request resources, such as new databases, or VMs with SQL Server installed. With vCloud Automation Center however, VMware do provide this functionality, yet it is not provided in the box and still relies on the customer to create and build the relevant services from scratch.

Once the services have been deployed, the DBA can access them directly via RDP, remotely from SQL Server Management Studio, or, for more IT admin control, IT can provide a rich self-service experience for the DBAs to access their SQL Server machines. Compare this with VMware, whom provide a number of different self-service experiences, all of which provide overly complex experiences for what needs to be a streamlined, simple process to provide access to the appropriate resources, VMs and SQL Servers.

Summary

Specifically when comparing Microsoft and VMware's offerings for virtualizing, managing and maintaining a SQL Server infrastructure, Microsoft delivers considerably more capability across virtualization, hybrid and management capabilities than those offered by VMware. Microsoft's deep relationships between its Windows Server, System Center, Microsoft Azure and SQL Server product groups, ensure that customers benefit from a battle-tested, deeply integrated set of solutions that enable them to get the most from their investments in the platform, and ensure they deliver the best SQL Server service within their IT infrastructures.