Microsoft

Best Practices for Virtualizing and Managing SQL Server

v1.0 – May 2013

Windows Server

Copyright Information

Table of Contents

Introduction

This guide provides high-level best practices and considerations for deploying and managing Microsoft SQL Server 2012 on a Microsoft virtualization infrastructure. The recommendations and guidance in this document aim to:

- Complement the architectural design of an organization's specific environment.
- Help organizations take advantage of the key platform features in SQL Server 2012 to deliver the highest levels of performance and availability.

Executive Summary

With the proliferation of applications in everyday use, enterprises have an increasing need for more instances of databases such as SQL Server. Most applications have their own set of database requirements, leading to the use of multiple versions of databases and to significant costs related to the databases and related hardware resources. More importantly, the hardware deployed for databases is not fully utilized or there is a demand to scale up hardware resources to meet peak utilization on certain databases. Therefore, it has become important to control the cost of databases and related hardware resources, and to optimize and scale the use of hardware resources in different scenarios to provide better flexibility and maintain service level agreements (SLAs). An ideal solution to this problem is virtualization.

Virtualization is fairly common now. Many organizations worldwide have moved beyond the nascent stage into being more advanced users of server virtualization, specifically. These organizations have gained benefits in the areas of costs, efficiency, operations, availability, agility, and resiliency.

Microsoft SQL Server 2012 and Windows Server 2012 provide a host of new features that can be used to effectively virtualize demanding database workloads that previously were not considered for virtualization. This guide explains these new features in the context of how organizations should consider virtualizing SQL Server 2012 on Windows Server 2012 and the benefits of managing this virtualized environment with Microsoft System Center 2012. Working together, these industry-leading products deliver an integrated platform with a low total cost of ownership (TCO) as well as mission-critical scale, performance, and availability. The platform also provides enhanced end-to-end security, management, and monitoring capabilities.

Further, many organizations now want to go a step beyond and adopt an IT infrastructure that is optimized for and ready to work in the cloud. They need an IT infrastructure that can seamlessly span from a private to a public cloud. To achieve this goal, organizations require a common virtualization platform that can increase performance and efficiency across the infrastructure. In this regard, Windows Server 2012 Hyper-V offers an ideal virtualization platform for SQL Server 2012.

Target Audience

This guide is intended for IT professionals and technical decision makers (TDMs), including IT consultants, architects, database administrators, and IT managers. They can use this guide to understand how to set up an environment for virtualizing SQL Server 2012 using an integrated virtualization platform built on some

of the latest Microsoft technologies, such as Windows Server 2012 Hyper-V and System Center 2012. Understanding key considerations and best practices can help TDMs effectively plan and deploy SQL Server 2012 virtualization using Windows Server 2012 Hyper-V. This guide serves the following purposes for these key roles:

- **Architects**: Understand how the entire virtualization environment will work as they design the architecture.
- **IT Managers**: Design processes to fit the overall virtualization environment so that costs are reduced and efficiency is increased as much as possible.
- **Database Administrators**: Understand how SQL Server 2012 can be set up and function in the virtual environment.

Scope

This guide focuses on providing an understanding of the key considerations for virtualizing SQL Server 2012 on a Windows Server 2012 host system, virtual machine, or other environment. At a broad level, the guide is divided into the following sections:

- **Fabric Configuration**: Covers the key requirements, features, and considerations for infrastructure that are necessary to set up the virtualization environment. This includes best practice considerations for physical hosts and compute requirements for processors, memory, storage, and networks.
- **Fabric/Host Resiliency**: Provides information related to Hyper-V host clustering and resiliency, and introduces features that enable resiliency on host systems, such as failover clustering and Cluster Shared Volume.
- **Virtual Machine Configuration for SQL Server and SQL Server Resiliency**: Highlights best practice considerations related to configuring virtual machines for SQL Server 2012. Provides information related to SQL Server resiliency in different scenarios, including a single cluster, clustering across multiple hosts, and Hyper-V Replica.
- **System Center Enhancements**: Provides an overview of how System Center 2012 SP1 supports deploying and managing SQL Server 2012 across the infrastructure (that is, on-premises and in the cloud).

Why Virtualize SQL Server?

In the past, database workloads were the exception for virtualization and were not viewed as prime candidates for consolidation because of their higher performance and scalability requirements.[1] Previous virtualization solutions had limitations that hindered their ability to meet the high performance and massive scalability requirements for workloads like SQL Server 2012.

With server virtualization today, organizations are able to lower IT capital and operational costs and attain greater IT efficiency with improved application provisioning, maintenance, availability, and backup/recovery processes. With recent technological advances, complex database workloads can be more readily consolidated using virtualization. The improvements in Windows Server 2012 Hyper-V and its super-sized virtual machine capabilities have eliminated most of the limits related to performance in a virtualized environment. Windows Server 2012 Hyper-V provides better consolidation of workloads that are traditionally more complex and that tend to saturate resources and contend for other system resources and storage. Such complex workloads include Online Transaction Processing/Analysis (OLTP/OLTA), Data Warehousing (DW), and Business Intelligence (BI).

Recently, the Enterprise Strategy Group surveyed organizations on their use of tier-2 database applications (such as Oracle Standard and Microsoft SQL) deployed on virtual machines in a production environment.[2] Of the 440 respondents, 50 percent had already deployed tier-2 database applications on production-based virtual machines, with 25 percent planning to in the near future (Figure 1).

Figure 1: Survey responses regarding tier-2 database applications

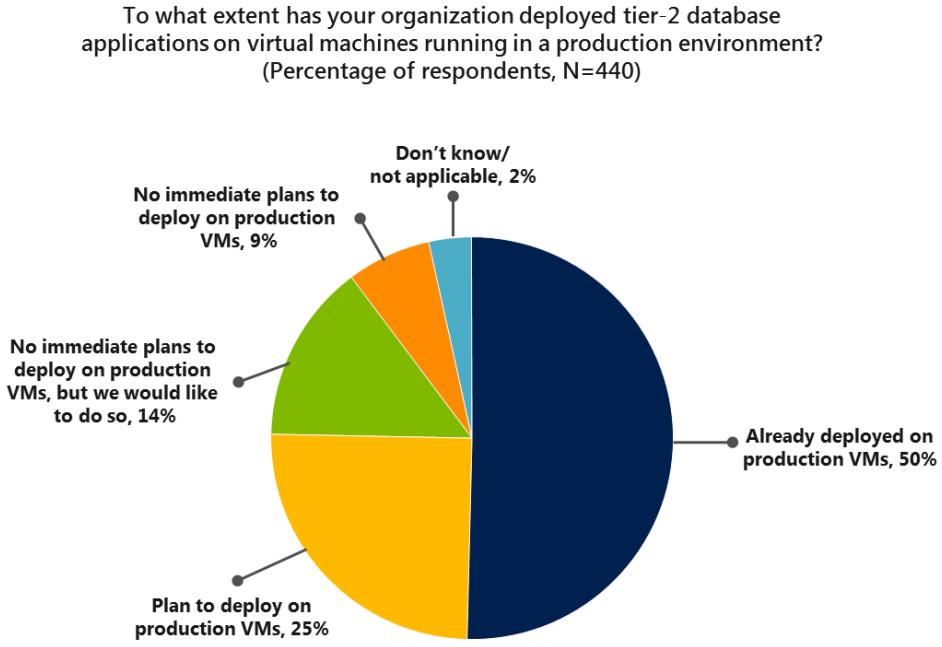

Why Microsoft Virtualization and Management?

Organizations today want the ability to consistently and coherently develop, deploy, and manage their services and applications across on-premises and cloud environments. Microsoft offers a consistent and integrated platform that spans from on-premises to cloud environments. This platform is based on key Microsoft technologies, including Windows Server 2012 Hyper-V, System Center 2012, Windows Azure, and Microsoft Visual Studio 2012.

Windows Server 2012 Hyper-V is an optimal virtualization platform that can be used for deploying demanding, mission-critical production applications, including SQL Server 2012 database workloads. With Hyper-V, Microsoft has become one of the leading vendors in virtualization technology.[3] This virtualization platform, based on new technologies from Microsoft, offers many features and improvements, including improved scale and performance, a hypervisor in the box, and enterprise features at no additional cost.

New and enhanced capabilities in Windows Server 2012 Hyper-V can help customers reduce costs, while at the same time providing increased agility and flexibility. Along with SQL Server 2012, Windows Server 2012 proves to be one of the best platforms for mission-critical workloads.[4] The platform provides:

- **Performance and Scalability**: Because SQL Server can use an operating system's maximum processor and memory capacity, Windows Server 2012 provides support for up to 640 logical processors (cores) over 64 sockets and up to 4 TB of memory (RAM). This enables organizations to meet peak demand by massively scaling up the most mission-critical applications built on SQL Server. The NIC Teaming feature in Windows Server 2012 helps improve the reliability of the networking system and ensures high availability of the network. It also improves the network throughput for SQL Server data. Windows Server 2012 features improvements in Server Message Block (SMB), including the ability to store SQL Server database files on remote file shares. It also offers multiple deployment options for database server storage.
- **Availability**: Windows Server 2012 supports running SQL Server 2012 on Windows Server Core. This helps to reduce the surface area for attack and reduces the need to apply patches at the operating system level. With the Cluster Aware Update feature of Windows Server 2012, maintenance of SQL Server cluster nodes can be automated, reducing downtime and improving reliability. With dynamic quorum management and the SQL Server AlwaysOn features, SQL Server clusters can achieve extremely high availability.

By combining Windows Server 2012 with System Center 2012, organizations can comprehensively manage demanding applications (such as SQL Server 2012 workloads) as well as the infrastructure—including physical and virtual resources—in an integrated and unified manner.[5] The key benefits of this integrated virtualization and management platform by Microsoft include the following:[4]

- **Better Scalability**: Higher capacity vCPUs (up to 64), memory (up to 1 TB), and virtual machine density (up to 8,000 per cluster).

- **Better Performance**: Hyper-V support on Non-Uniform Memory Access (NUMA) and Fibre Channel.
- **Better Availability**: Faster and simultaneous live migration and dynamic quorum support in SQL Server AlwaysOn cluster.
- **Better Manageability**: Same management tool (System Center) for SQL Server virtual machines in both private and public clouds.

Fabric Configuration

With Windows Server 2012 Hyper-V, organizations can take better advantage of the cost savings of virtualization and make the best use of server hardware investments by consolidating workloads as separate virtual machines. Windows Server 2012 provides a number of compelling capabilities to help organizations gain greater scalability and build reliable virtualized infrastructure for their mission-critical workloads like SQL Server 2012. This section covers many of the new enhancements in Windows Server 2012 Hyper-V that can help organizations build scalable, high-performance virtualized infrastructures.

Hardware Considerations

The hardware requirements of Windows Server 2012 Hyper-V help to ensure that it is installed correctly and makes optimum use of virtualization technology. Hyper-V requires a 64-bit processor that includes hardware-assisted virtualization and hardware-enforced Data Execution Prevention (DEP).

Hardware-assisted virtualization is available in processors that include a virtualization option to enable the full virtualization of host machines. The Windows Server 2012 Hyper-V role supports hardware-assisted virtualization processers from the Intel VT and AMD-V processor families. Using this feature, Hyper-V puts a layer between the processors enabled with hardware-assisted virtualization and the host operating system. This facilitates interaction between guest operating systems and the underlying hardware via the host or main operating system for better performance and control over hardware resources (Figure 2).

Figure 2: Full virtualization with Hyper-V

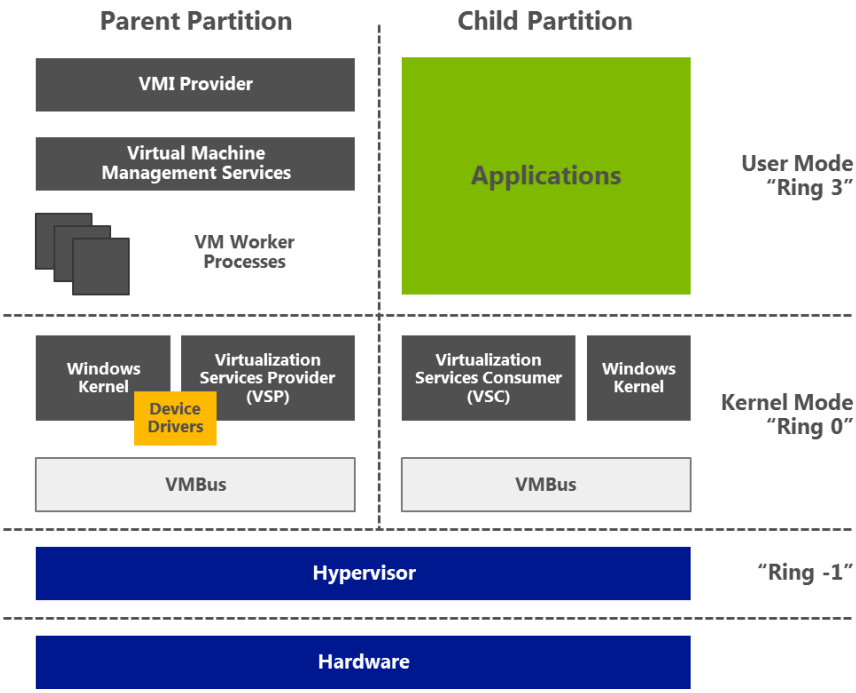

Hardware-enforced Data Execution Prevention must be available and enabled. Specifically, you must enable Intel XD bit (*execute disable bit*) or AMD NX bit (*no execute bit*).

The estimated minimum system requirements for Windows Server 2012 are as follows:[6]

- **Processor**: Minimum of 1.4 GHz 64-bit processor
- **Memory**: Minimum of 512 MB
- **Disk**: Minimum of 32 GB

These are minimum requirements only. The actual requirements will vary based on the system configuration used to create a virtualization environment with Windows Server 2012 Hyper-V and the applications and features installed. Therefore, we recommended carefully considering the intended SQL Server 2012 workloads and their requirements when planning for hardware resources.

For an optimal experience and better performance and stability of Windows Server 2012, we recommended using hardware and software that are compatible and carry the *Certified for Windows Server 2012* logo. Windows Server 2012 is compatible with most common hardware and software and has a large list of items from multiple manufacturers that are part of the Microsoft logo testing programs. The Windows Server Catalog lists thousands of hardware and software items compatible with Windows Server 2012. It is important to select the proper hardware to meet your expected performance and power goals because hardware bottlenecks limit the effectiveness of software tuning.

Windows Server 2012 provides different deployment options, including Server with a GUI and Server Core Installation. The *Server with a GUI* option is the Windows Server 2012 equivalent of the full installation option available in Windows Server 2008 R2. The *Server Core Installation* option reduces the space required on disk, the potential attack surface, and especially the requirements for servicing and restarting the server.[7]

Best Practices and Recommendations

Use the *Server Core Installation* option for setting up a SQL server virtualization environment. This helps to reduce the space required on disk and the potential attack surface. It also helps to better separate management and deployment activities, reduce the overhead of updating patches, and minimize the requirements for servicing and restarting the server.

Scalability Maximums of Windows Server 2012 Hyper-V

Windows Server 2012 Hyper-V provides significant scalability improvements over Windows Server 2008 R2 Hyper-V. Hyper-V in Windows Server 2012 greatly expands support for the number of host processors and memory for virtualization—up to 320 logical processors and 4 TB of physical memory, respectively. In addition, Hyper-V includes support for up to 64 virtual processors and 1 TB of memory per virtual machine, a new VHDX virtual hard disk (VHD) format with a larger disk capacity of up to 64 TB, and additional resiliency and alignment benefits. These features help to ensure that the virtualization infrastructure is compatible with the largest scale-up servers and can support the configuration of large, high-performance virtual machines to handle workloads that might need to scale up significantly.

Table 1 highlights additional improvements by comparing the resources supported by Hyper-V in Windows Server 2012 to those supported in Windows Server 2008 R2:[8,9]

Table 1: Resources available across versions of Windows Server

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

Significant improvements also have been made within Windows Server 2012 Hyper-V to support increased cluster size and a higher number of active virtual machines per host. Windows Server 2012 Hyper-V supports up to 8,000 virtual machines on a 64-node failover cluster. This is eight times and four times, respectively, the support provided by the previous version of Windows Server (that is, Windows Server 2008 R2).[10] In addition, more advanced performance features such as in-guest NUMA are supported by Windows Server 2012 Hyper-V virtual machines. Providing these enhancements helps to ensure that customers can achieve the highest levels of scalability, performance, and density for their mission-critical workloads.

Microsoft Assessment and Planning Toolkit

IT infrastructure for server virtualization requires proper planning, which includes gathering details related to the hardware that resides in current environments. The Microsoft Assessment and Planning Toolkit provides server utilization data for Hyper-V server virtualization planning, identifies server placements, and performs virtualization candidate assessments, including return on investment (ROI) analysis for server consolidation with Hyper-V.

Compute Considerations

Organizations need virtualization technology that can support the massive scalability requirements of demanding SQL Server workloads, such as Online Transaction Processing/Analysis (OLTP/OLTA), Data Warehousing (DW), and Business Intelligence (BI). One of the key requirements to virtualizing such workloads is to have a large amount of processing and memory power. Therefore, when planning to virtualize demanding, mission-critical, and high-performance workloads, you must properly plan for these compute resources.

Logical Processors on Hardware

Logical processors are representations or abstractions of a processor's physical cores themselves or of the number of threads that can be handled by a single physical core of a processor. Windows Server 2012 can run on host servers supporting up to 320 logical processors. With Windows Server 2012, there is no enforced limit on the virtual processor to logical processor (VP:LP) ratio. Users can have as many virtual processors associated with a logical processor as the hardware will allow. However, it is better to test the VP:LP ratio compatibility of the workload that needs to be virtualized to a level where the performance is not adversely affected.

Hyper-V also benefits from larger processor caches, especially for loads that have a large working set in memory and in virtual machine configurations where the VP:LP ratio is high.[11]

Customers can use processors that support **Second Level Address Translation (SLAT)** technologies (that is, SLAT-based processors). SLAT technologies add a second level of paging functionality under the paging tables of x86/x64 processors. They provide an indirection layer that maps virtual machine memory addresses to physical memory addresses, which reduces load on the hypervisor for address translation (Figure 3).

Figure 3: Virtual memory and SLAT

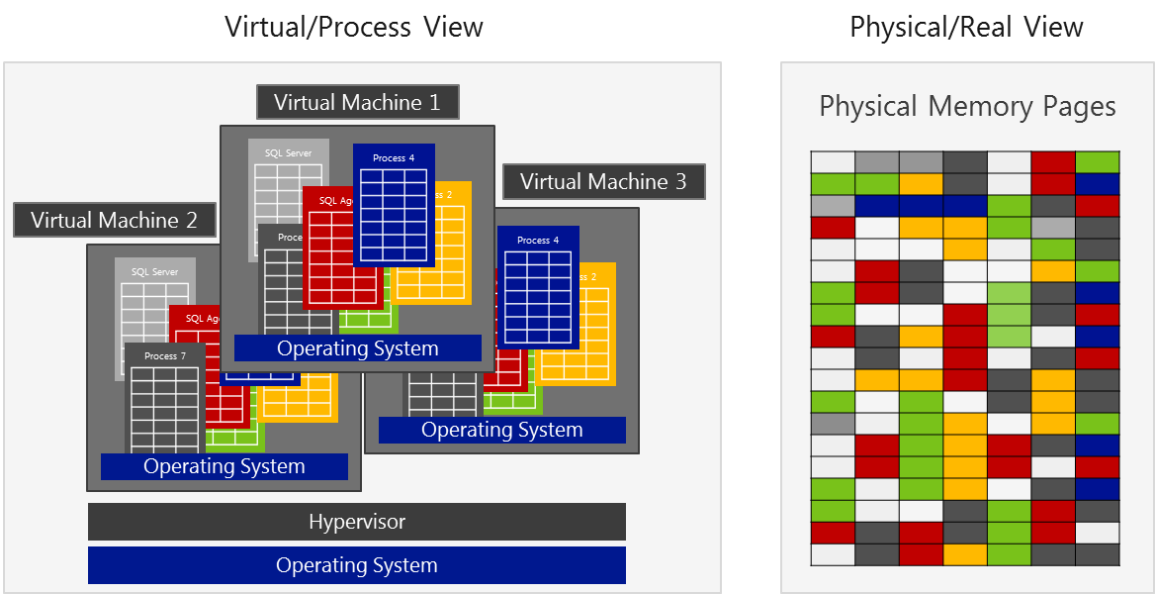

SLAT technologies also help to reduce CPU and memory overhead, thereby allowing more virtual machines to be run concurrently on a single Hyper-V machine. The Intel SLAT technology is known as *Extended Page Tables* (EPT); the AMD SLAT technology is known as *Rapid Virtualization Indexing* (RVI), formerly *Nested Paging Tables* (NPT).

Best Practices and Recommendations

For optimal performance of demanding workloads like SQL Server, run Windows Server 2012 Hyper-V on SLAT-capable processors/hardware. This offers the additional benefits of improved performance, more virtual machine density per host machine, and reduced overhead.

Virtual Processor or Virtual CPU

A virtual processor or a virtual CPU (vCPU) is a representation of the physical core of a processor or the threads/logical processors in the core. Virtual processors assigned to a virtual machine define its processing power. Hyper-V supports configuring virtual machines with more than one virtual processor from multiple physical or logical processors. In other words, one virtual machine can be configured to use multiple physical processor cores at the same time, thereby increasing performance. Such virtual machines are called Symmetric Multi-Processing (SMP) virtual machines. With SMP functionality, applications can fully benefit from multi-threading functionality while running in a virtualized environment.

As previously discussed, Windows Server 2012 Hyper-V supports virtual machines with up to 64 virtual processors and 1 TB of memory. With increased support for 320 logical processors on a host machine/ hardware, Hyper-V in Windows Server 2012 can now support up to 2,048 virtual processors per host.

Note

Unlike the earlier version of Windows Server, with Windows Server 2012, there is no VP:LP ratio imposed by Hyper-V.

Windows Server 2012 Hyper-V also provides the Weights and Reserves feature (Figure 4). Weights are assigned to a virtual processor to grant it a larger or smaller share of CPU cycles than the average cycle share. Reserves are set for a virtual processor to ensure that it gets at least a specified percentage of the total possible CPU usage of a virtual machine when there is contention for CPU resources. Simply put, if there is higher demand for CPU than is physically available, Hyper-V ensures that a virtual machine needing CPU resources gets at least its CPU reserve when there is contention.[12] This feature is especially beneficial for system administrators who want to prioritize specific virtual machines depending on the load they have or need to handle.

Figure 4: Weights and reserves in Windows Server 2012

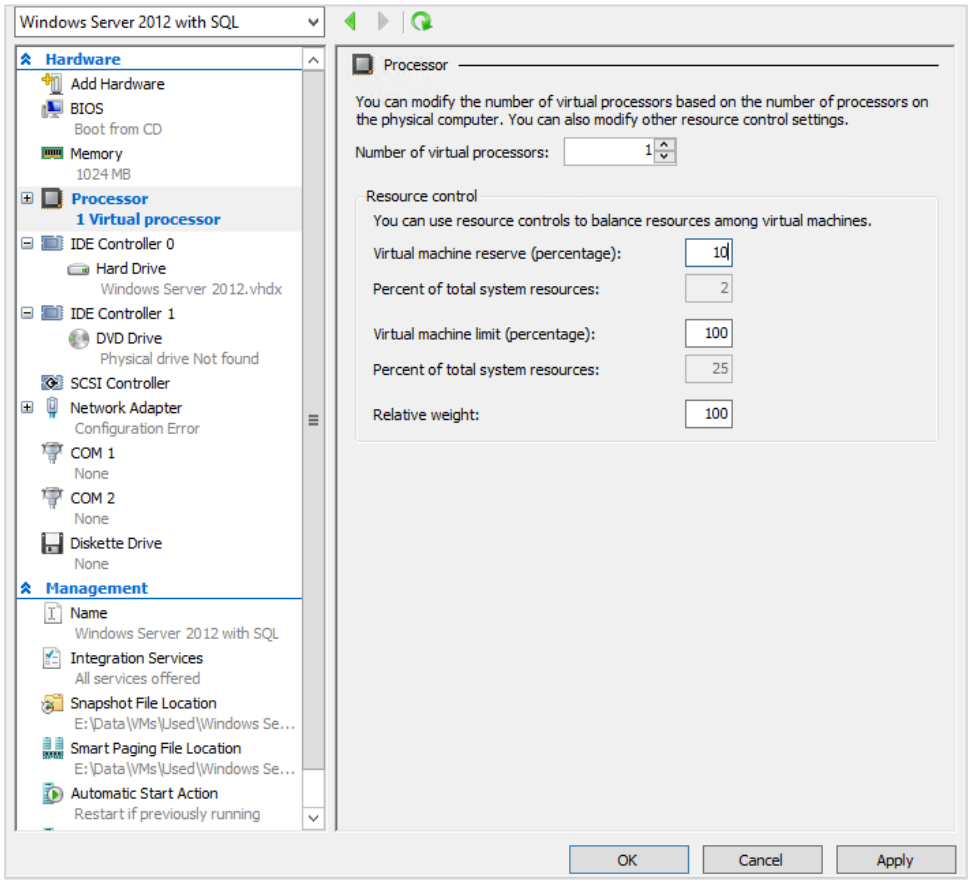

Best Practices and Recommendations

The Weights and Reserves feature, when used properly, can be a great tuning mechanism. If CPU resources are overcommitted, you can set weights and reserves to optimize the way these resources are used. You can prioritize or deprioritize specific virtual machines based on the intensity of loads they bear (that is, high-intensity and low-intensity loads).

Non-Uniform Memory Access – Host Perspective

In single system bus architecture, all processors fetch memory from a single pool, and all requests for memory are sent using a single system bus. One problem with this architecture is that as the speed and number of processors increase, it becomes difficult for the system to handle a large number of memory requests. This leads to issues such as memory latency and scalability limitations. While one solution for such issues is to have larger cache size, this helps only to a certain extent. The issues related to memory access can be best resolved with NUMA.[13]

NUMA is a memory design architecture that provides significant advantages over the single system bus architecture and provides a scalable solution to memory access problems. In a NUMA-supported

operating system, CPUs are arranged in smaller systems called *nodes* (Figure 5). Each node has its own processors and memory, and is connected to the larger system through a cache-coherent interconnect bus.[14]

Figure 5: NUMA node (processor and memory grouped together)

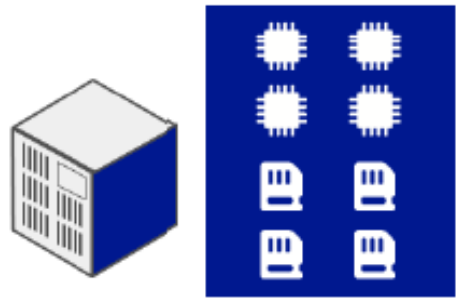

Multiple NUMA nodes can exist in a host system (Figure 6). In the context of multiple nodes:

- Local memory is attached directly to the processor (grouped into a node).
- Remote memory is local to another processor in the system (another node).

This grouping into nodes reduces the time required by a processor to access memory (locally located), as the processor can access local memory faster than remote memory.[15]

Figure 6: Multiple NUMA nodes on a single host

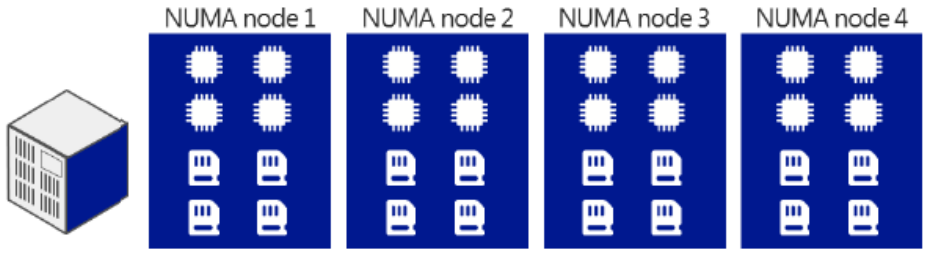

Root/Host Reserve

Root reserve or host reserve is the amount of memory that is reserved for the root partition and is guaranteed to be available to the root partition. It is not allocated to any of the virtual machines running in the child partition. Hyper-V automatically calculates root reserve based on the physical memory available on the host system and system architecture.[16]

Best Practices and Recommendations

The root partition must have sufficient memory to provide services such as input/output (I/O) virtualization, virtual machine snapshot, and management to support the child partitions. Hyper-V calculates an amount of memory (known as the root reserve), which is guaranteed to be available to the root partition. This memory is never assigned to virtual machines. Root reserve is calculated automatically, based on the host's physical memory and system architecture.

Page File Guidance

When a machine runs low on memory and needs more immediately, the operating system uses hard disk space to supplement system RAM through a procedure called *paging*. Too much paging degrades overall system performance. However, you can optimize paging by using the following best practices and recommendations for page file placement.

Best Practices and Recommendations

Let Windows Server 2012 handle the page file sizing. It is well optimized in this release.

Isolate the page file on its own storage devices, or at least make sure it does not share the same storage devices as other frequently accessed files. For example, place the page file and operating system files on separate physical disk drives.

Place the page file on a drive that is not fault-tolerant. Note that if the disk fails, a system crash is likely to occur. If you place the page file on a fault-tolerant drive, remember that fault-tolerant systems are often slower to write data because they do so to multiple locations.

Use multiple disks or a disk array if you need additional disk bandwidth for paging. Do not place multiple page files on different partitions of the same physical disk drive.

The following additional best practices and recommendations should be considered while planning and managing host compute (CPU and memory) resources.[17]

Best Practices and Recommendations

While performing capacity planning for virtualizing workloads, always count the number of cores required and not the number of logical processors/threads required.

When you are planning how to use the host server's memory, it is important to consider the virtualization-related overhead. Whether you choose to use NUMA or Dynamic Memory, both have some overhead related to memory management in the virtualized environment. There may be scenarios when using NUMA or Dynamic Memory may not be the best option. For example, if an application or workload works best on a fixed amount of memory, it is good to allocate to that virtual machine the exact amount of memory required. Therefore, we recommend that you properly plan the memory requirements for running SQL Server 2012 workloads on Windows Server 2012, and choose the memory management technology that is most beneficial for the workload.

To mitigate performance impacts on SQL Server, place the Windows page file on high-performance storage. This is helpful when, in the event of low memory, Windows needs to page out the working set of SQL Server.

Storage Considerations

Storage configuration is one of the critical components to any successful database deployment. Database servers tend to be heavily I/O bound due to rigorous database read and write activity and transaction log processing. As a result, improper configuration of I/O subsystems can lead to poor performance and operation of SQL Server systems. In a virtualized SQL Server deployment, Windows Server 2012 storage capabilities are designed to provide enhanced reliability, availability, and performance for widely distributed database applications.

Storage Options for Hyper-V Virtual Machine

Storage virtualization helps administrators perform backup, archiving, and recovery tasks more easily and quickly by reducing the complexity of storage devices and the time required to manage them. Windows Server 2012 introduces a class of sophisticated storage virtualization enhancements that can be easily implemented to develop resilient databases. These enhancements use two new concepts: Storage Spaces and Storage Pools.

Storage Spaces

With the Storage Spaces technology, you can achieve a desired level of resiliency through automatic or controlled allocation of heterogeneous storage media presented as one logical entity. Storage Spaces shields the physical disks and presents selected storage capacity as pools, known as *storage pools*, in which a virtual disk, known as a *storage space*, can be created. Storage Spaces supports two optional resiliency modes: mirroring and parity. These provide per-pool support for disks that are reserved for replacing failed disks (hot spares), background scrubbing, and intelligent error correction. In case 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.

The Storage Spaces technology is fully integrated with failover clustering to enable continuously available service deployments. One or more storage pools can be clustered across multiple nodes within a single cluster. Storage Spaces supports thin provisioning to allow organizations to easily share storage capacity among multiple unrelated data sets, thereby maximizing capacity use. Fully scriptable management is enabled through the Windows Storage Management API, Windows Management Instrumentation (WMI), and Windows PowerShell. Storage Spaces also can be managed through the File and Storage Services role in Server Manager. Finally, Storage Spaces provides notifications when the amount of available capacity in a storage pool hits a configurable threshold.

Storage Pools

Storage pools are a collection of disks used for storing replicas, shadow copies, and transfer logs and are the fundamental building blocks for Storage Spaces (Figure 7). In Windows Server 2012, storage pools are a collection of physical disks grouped together into one or more containers. This allows for storage aggregation, flexible capacity expansion, and delegated administration of storage. Windows Server 2012 maps a storage pool by combining a group of hard disks and/or solid-state drives (SSDs). By simply adding additional drives, storage pools are dynamically expanded to handle the growing size of data. Thinly provisioned virtual disks can be provisioned from the available capacity. Thin provisioning helps in reserving the actual capacity by reclaiming capacity on the space whenever files are deleted or no longer in use.

Figure 7: Conceptual deployment model for storage spaces and storage pools

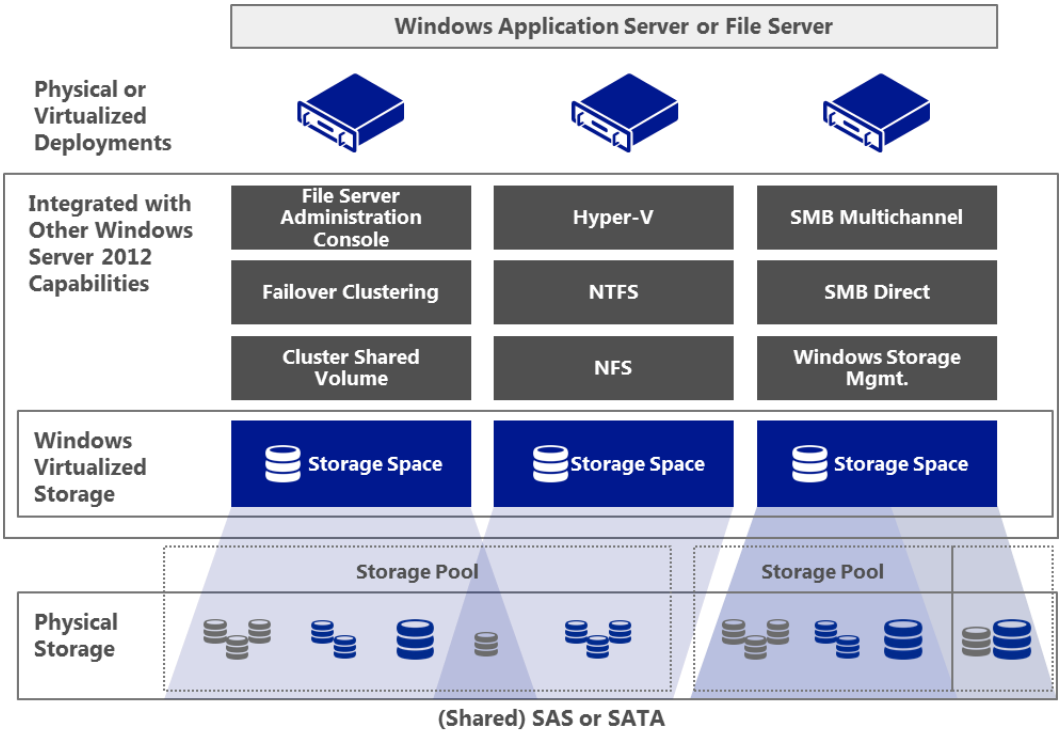

Types of Storage Spaces

There are three key types of storage spaces: simple/striped spaces, mirror spaces, and parity spaces. Each is discussed in more detail below.[18]

Simple Spaces/Striped Spaces: Simple storage spaces are used for storing temporary data because they are non-resilient to disk failures. Striping is the process of writing data across multiple disks to reduce access and response times. Logical blocks of data with a defined size are laid out in a sequential circular manner across multiple disks. This helps in balancing the storage load across all physical drives. Striping provides the overall best performance in terms of reads and writes but, as noted, provides no resiliency.

In Figure 8, there are four disks, and 1 MB of data needs to be written to these disks. In this case, there are two options for writing data to the disks: Either write all of the data to a single disk and access it from there, or write 256 KB to each of the four disks simultaneously. The second option results in a quadruple decrease in write times. The greater the number of disks Storage Spaces can stripe across, the better the performance will be.

Figure 8: Striped storage space across four disks

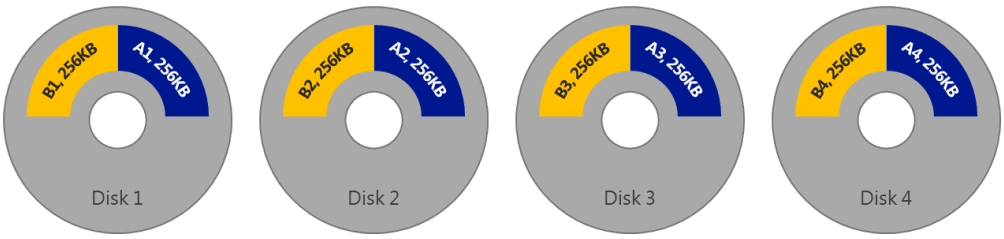

Striped storage spaces can be used for the following:

- Delivering the overall best performance in terms of reads and writes.
- Balancing the overall storage load across all physical drives.
- Backing up disks to increase backup throughput or to distribute the use of space across different disks.

Mirror Spaces: This data layout process uses the concept of mirroring to create copies of data on multiple physical disks. A logical virtual disk is created by combining two or more sets of mirrored disks. Mirror storage spaces are resilient in nature because in the event of failure, if one copy is lost, the other is still available. To make them resilient from disk failures, mirror spaces are configured to at least one (two-way mirror) or two (three-way mirror) concurrent physical disks.

In Figure 9, 512 KB of data needs to be written to the storage space. For the first stripe of data (A1), Storage Spaces writes 256 KB of data to the first column, which is written in duplicate to the first two disks. For the second stripe of data (A2), Storage Spaces writes 256 KB of data to the second column, which is written in duplicate to the next two disks. The column-to-disk correlation of a two-way mirror is 1:2, while for a three-way mirror, the correlation is 1:3. Reads on mirror spaces are very fast because they are done from either of the two copies of data. If disks 1 and 3 are busy servicing another request, the needed data can be read from disks 2 and 4.

Figure 9: Mirror storage space across four disks

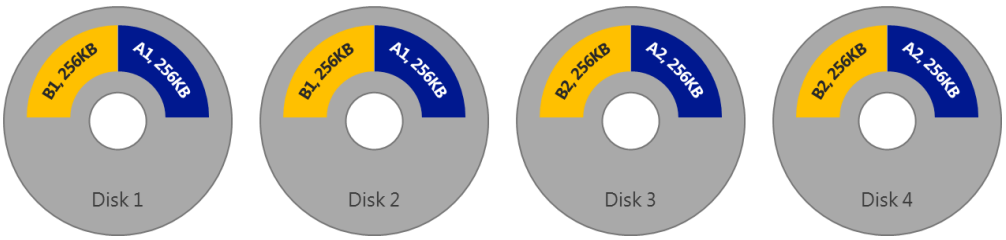

Mirror storage spaces are used for the following:

- Enabling faster reads on data.
- Increasing backup reliability to prevent loss from a malfunctioning backup device.

Parity Spaces: Parity storage spaces store parity-bit information that helps in reconstructing data from a failed disk. This can be useful in providing data recovery capabilities. Storage Spaces uses rotating parity that stores data and parity information by rotating from stripe to stripe across different disks. Parity spaces tend to have lower write performance than mirror spaces because each parity block takes time in updating itself to the corresponding modified data block. Parity is more cost efficient than mirroring because it requires only one additional disk per virtual disk, instead of double or triple the total number of disks in an array.

In Figure 10, for the first stripe of data, 768 KB is written across disks 1 through 3 (A1, A2, A3), while the corresponding parity bit (AP) is placed on disk 4. For the second stripe of data, Storage Spaces writes the data on disks 1, 2, and 4, thereby rotating the parity to disk 3 (BP). Because parity is striped across all disks, it provides good read performance and resiliency to single disk failure.

Figure 10: Parity storage space across four disks

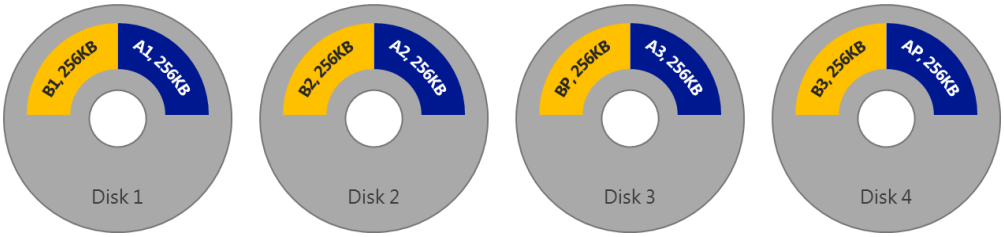

Parity storage spaces are used for the following:

- Providing data recovery of failed disks.
- Offering efficient capacity utilization.
- Delivering faster read operations.
- Providing bulk backups by writing data in large sequential append blocks.

Best Practices and Recommendations

When you deploy storage spaces, consider the following:

- Use a clustered storage pool made up of at least three SAS-connected physical disk drives, each with at least 4 GB capacity. The physical disks used for a clustered storage pool must be dedicated to that pool. Startup disks should not be added.
- Use either simple/striped or mirror storage spaces. Note that parity storage spaces are not supported for use with virtualized SQL Server 2012.

The graphs in Figure 11 show the performance scaling of a simple storage space with up to 32 disks, which resulted in a random read 1.4 million IOPS and 10.9 GB/sec of sequential throughput.[19]

Figure 11: Performance scaling of a simple storage space

The Enterprise Strategy Group (ESG) Lab tested the Storage Spaces technology in an OLTP online brokerage application to simulate the activity of thousands of SQL Server 2012 users.[20] The goal was to demonstrate the performance and scalability of Storage Spaces.

Each of eight virtual machines was assigned a volume consisting of a 2 TB database and a 90 GB log. The workload used during ESG Lab testing was designed to emulate the database activity of users in a typical online brokerage firm as they generated trades, performed account inquiries, and did market research. The workload was composed of 10 transaction types with a defined ratio of execution. Four of the transactions performed database updates; the rest were read-only. A database of 3,000 customers was configured within each virtual machine, with a database scaling goal of 24,000 customers. The workload generated a high level of I/O activity with small access sizes and spent a lot of execution time at the operating system kernel level. These characteristics, combined with a large, cache-resident working set, created a workload that was well suited for evaluating the performance, efficiency, and scalability of Storage Spaces.

The average response time for the 10 transaction types was monitored as the scale of the databases and number of concurrent users increased within the virtual machines (Figure 12). Performance was also monitored through the number of disk transfers/sec (IOPS) and SQL batch requests/second.

Figure 12: SQL Server 2012 scalability with Storage Spaces

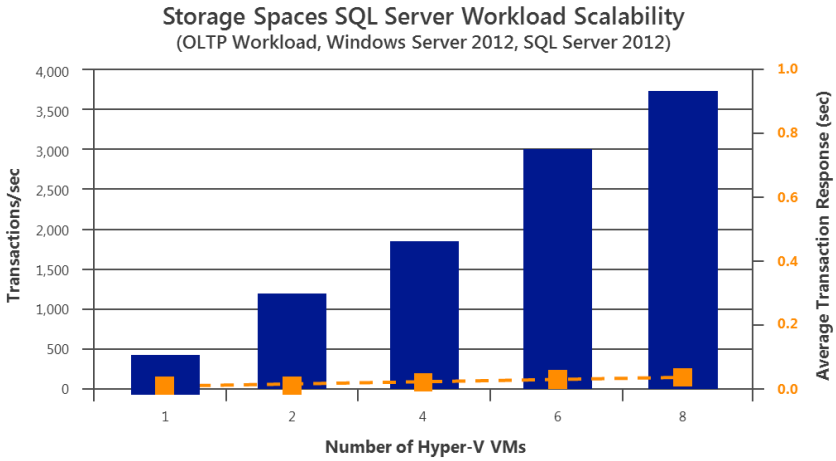

Storage Protocols and Additional Features

Various storage protocols can help in virtualizing workloads to connect easily and reliably to existing storage arrays. These storage protocols include a vast number of storage feature enhancements that increase administrative flexibility, efficiency, and control by centralizing management of storage volumes. Apart from storage protocols, Windows Server 2012 allows efficient data movement using intelligent storage arrays and enables rapid provisioning and migration of virtual machines. Some of these storage protocols and features are described below.

Server Message Block 3.0

The Server Message Block (SMB) protocol is a network file sharing protocol that allows applications to read, create, update, and access files or other resources at a remote server. The SMB protocol can be used on top of its TCP/IP protocol or other network protocols. Windows Server 2012 introduces the new 3.0

version of the SMB protocol that greatly enhances the reliability, availability, manageability, and performance of file servers. SMB 3.0 also allows you to create a failover cluster without shared storage or expensive storage area networks (SANs).

Hyper-V over SMB

By enabling Hyper-V to use SMB file shares, performance can be greatly enhanced with easy and inexpensive deployments of virtual storage. Hyper-V over SMB can be used to keep virtual storage (.vhd and .vhdx files) on a remote file server rather than requiring the Hyper-V host to manage the storage for its many virtual machines. This allows Hyper-V hosts to provide compute resources with many processors and RAM and to connect storage to file servers as virtual storage. To enable Hyper-V over SMB host requires:

- One or more computers running Windows Server 2012 with the Hyper-V and File and Storage Services roles installed.
- A common Active Directory infrastructure. (The servers running Active Directory Domain Services do not have to run Windows Server 2012.)

Failover clustering on the Hyper-V side, the File and Storage Services side, or both is optional.

Hyper-V over SMB supports a variety of flexible configurations that offer different levels of capabilities and availability. These configurations include Single-Node File Server, Dual-Node File Server, and Multi-Node File Server, as shown in the following figures.[21]

Single-Node File Server: In a Single-Node File Server, Hyper-V shares are used for VHD storage (Figure 13). File servers use standalone and local storage. This configuration provides flexibility for shared storage, as well as low costs for acquisition and operation. It does not provide continuous availability. Storage is not fault-tolerant, and Hyper-V virtual machines are not highly available.

Figure 13: Single-Node File Server

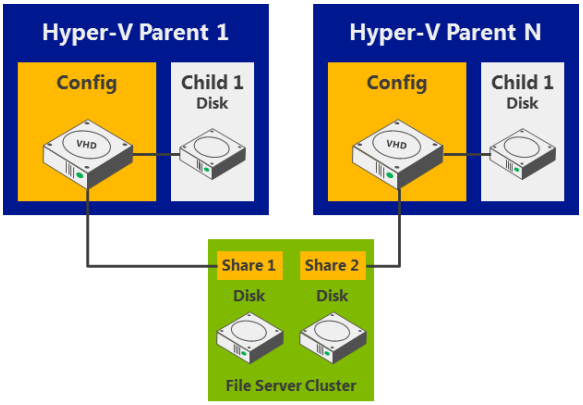

Dual-Node File Server: In a Dual-Node File Server, file servers can be clustered storage spaces, where shares are used for VHD storage (Figure 14). This configuration provides flexibility for shared storage, fault-tolerant storage, and low costs for acquisition and operation. It also offers continuous availability but with limited scalability.

Figure 14: Dual-Node File Server

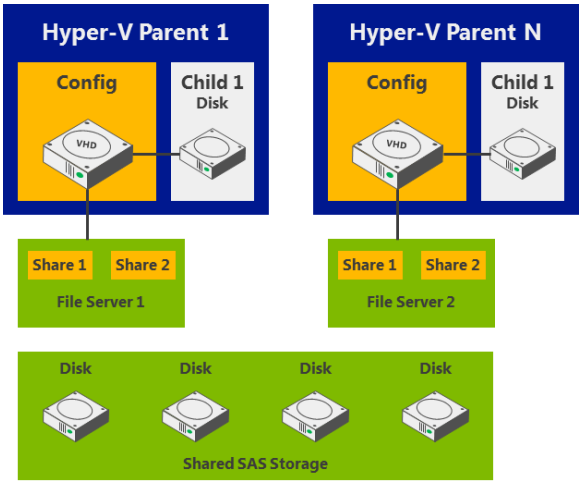

Multi-Node File Server: A Multi-Node File Server uses clustered Hyper-V file servers and storage spaces, where shares are used for VHD storage (Figure 15). This configuration provides flexibility for shared storage, fault-tolerant storage, and low costs for acquisition and operation. It also provides continuous availability, and Hyper-V virtual machines are highly available.

Figure 15: Multi-Node File Server

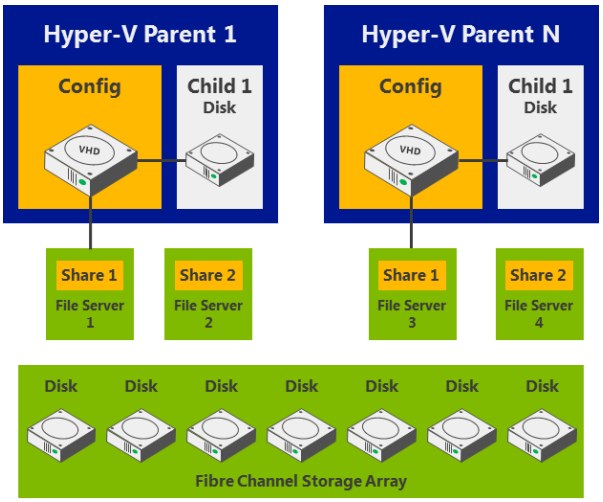

Table 2 compares the cost and availability/scalability of the three configurations for Hyper-V over SMB.

Table 2: Comparison of Hyper-V over SMB configurations

	Single-Node File Server	Dual-Node File Server	Multi-Node File Server
Cost	Lowest cost for shared storage	Low cost for continuously available shared storage	Higher cost, but still lower than connecting all Hyper-V hosts with Fibre Channel (FC)
Availability/ Scalability	Shares not continuously available	Limited scalability (up to a few hundred disks)	Highest scalability (up to thousands of disks)

SMB Multichannel

Both the SMB client and SMB server must support SMB 3.0 to take advantage of the SMB Multichannel functionality. SMB Multichannel increases the network performance and availability for file servers. SMB Multichannel allows file servers to use multiple network connections simultaneously. This increases throughput by transmitting more data using multiple connections for high-speed network adapters or multiple network adapters. When using multiple network connections at the same time, the clients can continue to work uninterrupted despite the loss of a network connection. SMB Multichannel automatically discovers the existence of multiple available network paths and dynamically adds connections as required.

Best Practices and Recommendations

If you use SMB storage with Hyper-V, use multiple network adaptors to take advantage of SMB Multichannel.

SMB Direct (SMB over RDMA)

Windows Server 2012 introduces SMB Direct, a feature that provides the ability to use Remote Direct Memory Access (RDMA) network interfaces for high throughput with low latency and CPU utilization. SMB Direct supports the use of network adapters that have RDMA capability. Network adapters with RDMA can function at full speed with very low latency, while using very little CPU. For workloads such as Hyper-V or SQL Server, this enables a remote file server to resemble local storage. SMB Direct is automatically configured by Windows Server 2012 and includes the following benefits:

- **Increased throughput**: Takes advantage of the full throughput of high-speed networks where the network adapters coordinate the transfer of large amounts of data at line speed.
- **Low latency**: Provides extremely fast responses to network requests and, as a result, makes remote file storage feel as if it is directly attached block storage.
- **Low CPU utilization**: Uses fewer CPU cycles when transferring data over the network, which leaves more power available to server applications.

By supporting mission-critical application workloads, the new SMB server and client cooperate to provide transparent failover to an alternative cluster node for all SMB operations for planned moves and unplanned failures. This results in reduced cost, improved high availability, and increased performance for workloads in a virtualized environment.

Best Practices and Recommendations

SMB Direct works with SMB Multichannel to transparently provide exceptional performance and failover resiliency when multiple RDMA links between clients and SMB file servers are detected. Also, because RDMA bypasses the kernel stack, it does not work with NIC Teaming, but does work with SMB Multichannel (because SMB Multichannel is enabled at the application layer).

Non-Microsoft file servers that implement the SMB 3.0 protocol from storage partners like EMC and NetApp can be used with Hyper-V.

Best Practices and Recommendations

Loopback configurations are not supported by Hyper-V when the file server is configured on the host where the virtual machines are running.

The ESG Lab tested the SMB 3.0 protocol by using an OLTP workload application to simulate the activity of SQL Server users (Figure 16). The goal was to demonstrate the performance, scalability, and efficiency of the SMB protocol, Hyper-V hypervisor, and SQL Server database engine on cost-effective commodity hardware.

A database of 3,000 customers was configured within each of eight SQL Server virtual machines, with a goal of achieving linear scalability for the number of transactions per second as the number of consolidated SQL Server virtual machines increased. The transactions per second and average response time were monitored as the number of customers and virtual machines increased.

Figure 16: Workload scalability with SMB 3.0 and SQL Server

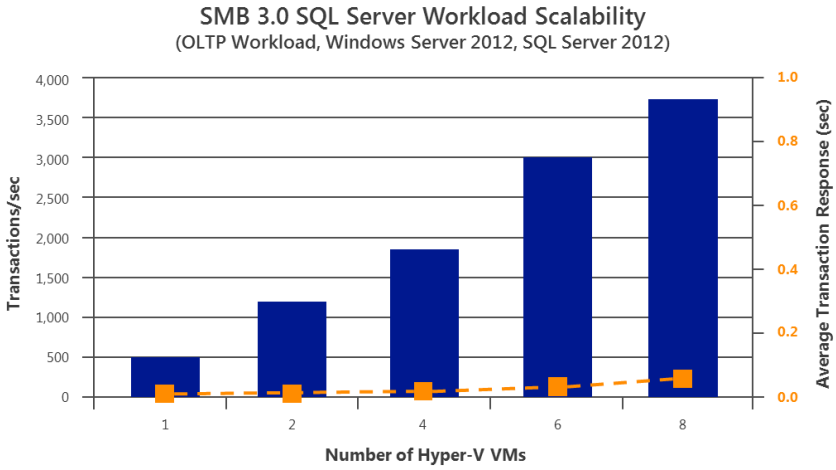

Internet SCSI

The Internet Small Computer System Interface (iSCSI) protocol is based on a storage networking standard that facilitates data transfers over the Internet and manages storage over long distances, all while enabling hosts (such as databases and web servers) to operate as if the disks were attached locally.

An iSCSI target is available as a built-in option in Windows Server 2012; it allows sharing block storage remotely by using the Ethernet network without any specialized hardware. It also provides support for diskless network boot capabilities and continuous availability configurations.

Fibre Channel

Fibre Channel (FC) is a data transmitting technology that enables server-to-storage connectivity at 16 GB and is well suited for connecting storage controllers and drives. Fibre Channel offers point-to-point, switched, and loop interfaces. It is designed to interoperate with SCSI, the Internet Protocol (IP), and other protocols. With the new 16 GB FC, a bi-directional throughput of 3,200 MB/sec can be delivered over 1 million IOPS. This enhancement supports deployments of densely virtualized servers, increases scalability,

and matches the performance of multicore processors and SSD-based storage infrastructure. 16 GB FC is backward compatible with 8/4 GB FC, allowing them to be seamlessly integrated into expansion segments of existing FC networks.

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) offers the benefits of using an Ethernet transport while retaining the advantages of the FC protocol and the ability to use FC storage arrays. This solution helps to reduce costs in several ways, including the elimination of dedicated FC switches and a reduction in cabling (which can be a significant cost in large data center environments). For higher performance and availability, FCoE provides direct connections to the FC host bus adapter (HBA) and SAN fabric from Hyper-V virtual machines.

Multipath I/O

Microsoft Multipath I/O (MPIO) is a framework provided by Microsoft for developing multipath solutions that contain hardware-specific information required to enhance connectivity for storage arrays. In other words, MPIO increases the availability of storage resources by providing support for using multiple data paths to a storage device. MPIO uses host-based software, called *device-specific modules* (DSMs), to provide this multipath support. MPIO is protocol-independent and can be used with FC, iSCSI, and Serial Attached SCSI (SAS) interfaces in Windows Server 2012. MPIO in Windows Server 2012 provides the following enhanced features:

- **PowerShell management and configuration**: MPIO can be configured using PowerShell as an alternative to MPCLAIM.exe.
- **Heterogeneous HBA usage with MPIO**: Heterogeneous HBA types now can be used together with non-boot virtual disks only.
- **Support for MPIO with multiport-SAS enclosures**: The use of MPIO with data volumes on a multiport-SAS enclosure is now supported.

An MPIO/multipath driver cannot work effectively until it discovers, enumerates, and configures into a logical group the different devices that the operating system sees through redundant adapters. Figure 17 shows that without any multipath driver, the same devices through different physical paths would appear as different devices, leaving room for data corruption.

Figure 17: The use of multipathing software to correctly identify paths and devices

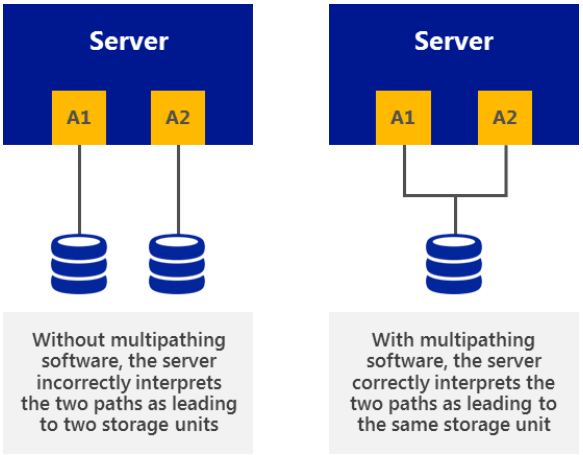

With MPIO, Windows Server 2012 efficiently manages 32 paths between storage devices and the Windows host operating system, and provides fault-tolerant connectivity to storage. Further, as more and more data is consolidated on SANs, the potential loss of access to storage resources is unacceptable. To mitigate this risk, high availability solutions like MPIO have become a requirement.

MPIO provides the logical facility for routing I/O over redundant hardware paths connecting servers to storage. These redundant hardware paths are composed of components such as cabling, HBAs, switches, storage controllers, and possibly even power. MPIO solutions logically manage these redundant connections so that I/O requests can be rerouted if a component along one path fails. The MPIO software supports the ability to balance I/O workload without administrator intervention. MPIO determines which paths to a device are in an active state and can be used for load balancing. Each vendor's load balancing policy setting is set in the DSM. (Individual policy settings may use any of several algorithms—such as Round Robin, Least Queue Depth, Weighted Path, and Least Blocks—or a vendor-unique algorithm.) This policy setting determines how I/O requests are actually routed.

Best Practices and Recommendations

To determine which DSM to use with existing storage, it is important to check with the storage array manufacturer. Multipath solutions are supported as long as a DSM is implemented in line with logo requirements for MPIO. Most multipath solutions for Windows use the MPIO architecture and a DSM provided by the storage array manufacturer. Use the Microsoft DSM provided in Windows Server only if it is also supported by the storage array manufacturer, in lieu of the manufacturer providing its own DSM.

A DSM from the storage array manufacturer may provide additional value beyond the implementation of the Microsoft DSM because the software typically provides auto-configuration, heuristics for specific storage arrays, statistical analysis, and integrated management. We recommend that you use the DSM provided by the storage array manufacturer to achieve optimal performance. This is because storage array manufacturers can make more advanced path decisions in their DSMs that are specific to their arrays, which may result in quicker path failover times.

Offloaded Data Transfer

Offloaded Data Transfer (ODX) in Windows Server 2012 allows rapid provisioning and migration of large files, such as database or video files and virtual machines, between existing storage arrays. By offloading the file transfer to the storage array, ODX minimizes delays, maximizes the use of array throughput, and reduces host resource usage, such as CPU and network consumption. Whenever a file is moved or copied, ODX automatically transfers and transparently offloads it without the need for administrative setup or intervention. In this way, ODX prevents the unnecessary steps required by traditional host-based file transfers.

In virtualized SQL Server 2012, ODX uses a token-based mechanism for reading and writing data within or between intelligent virtual storage database volumes (Figure 18). Instead of routing the data through the host, a small token is copied between the source and destination SQL Server virtual machines. The token serves as a point-in-time representation of the data. For example, when a database file is copied or a virtual SQL Server machine is migrated between storage locations, Windows Server 2012 copies the token representing the virtual machine file. This removes the need for the virtual machine server to copy the

underlying data between servers. In this way, ODX in Windows Server 2012 enables rapid provisioning and migration of virtualized database files and lets the database administrator quickly move large database instances directly between locations on the storage arrays—all while reducing host CPU and network resource consumption.

Figure 18: Offloaded Data Transfer in Windows Server 2012

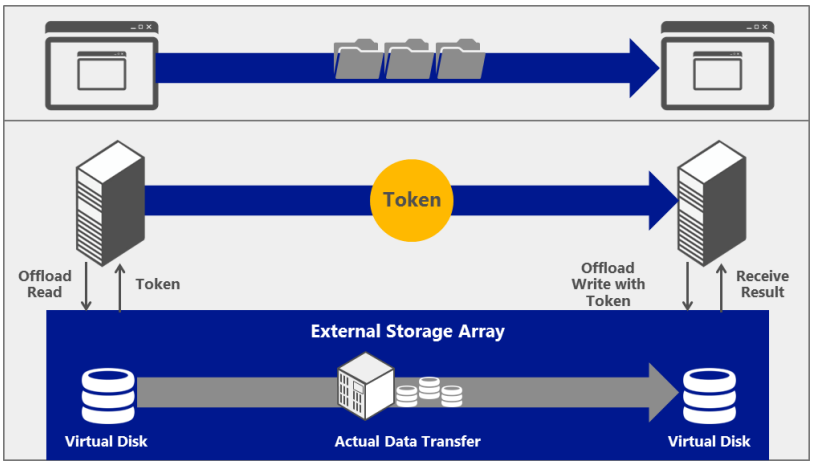

The ESG Lab tested the efficiency and functionality of Offloaded Data Transfer. Two servers were connected to an ODX-compliant Dell EqualLogic storage array. The storage array consisted of 12 SAS drives (600 GB each). A single RAID5 pool was created with two volumes: One contained a 75 GB virtual machine, and the other was empty. Using an intuitive wizard, the ESG Lab configured a virtual machine live migration from one server to another within a SAN. The lab specified the type of move, the server receiving the data, move options, and destination virtual machine options. It then transferred a virtual machine using the traditional non-ODX method and the new ODX method. The lab monitored network utilization and elapsed time for the transfer to complete in both test cases.

The results in Figure 19 show noticeable improvements using ODX. The ODX transfer took approximately 6.5 minutes for the virtual machine to completely migrate to the other server, and the average network bandwidth consumption was around 64 Kb/sec. Conversely, with non-ODX method, moving the 75 GB virtual machine over the network took approximately 52 minutes and consumed 4 Mb/sec of network bandwidth. The ODX method completed eight times faster than the non-ODX method, while consuming virtually no server CPU or network resources.

Figure 19: Faster SAN-attached virtual machine migrations with ODX

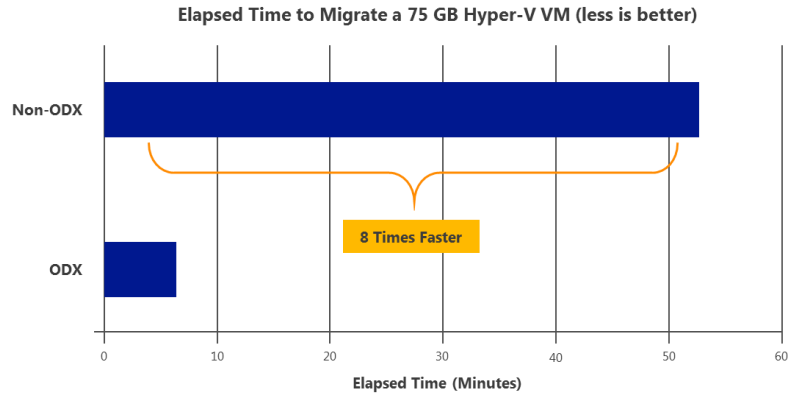

Best Practices and Recommendations

If you are using SAS or FC in all clustered servers, all elements of the storage stack should be identical. It is required that the MPIO and DSM software be identical. It is recommended that the mass storage device controllers (that is, the HBA, HBA drivers, and HBA firmware attached to cluster storage) be identical.[22]

If you are using iSCSI, each clustered server should have one or more network adapters or iSCSI HBAs that are dedicated to the cluster storage. The network being used for iSCSI should not be used for network communication. In all clustered servers, the network adapters being used to connect to the iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or higher. Network adapter teaming (also called *load balancing and failover*, or LBFO) is not supported for iSCSI. MPIO software can be used instead.

ODX is enabled by default, but check with your storage vendor for support, as upgraded firmware may be required.

Networking Considerations

Networking is one of the most important components in computing environments where database applications connect to databases. Typically, a database application uses a client/server architecture where the client portion accesses data from the database server through the network. Therefore, network performance and availability is important for mission-critical application data like that in SQL Server 2012.

Windows Server 2008 R2 introduced several networking-related features that help to reduce networking complexity while simplifying management tasks. Windows Server 2012 improves on this functionality in several ways, including through features such as NIC Teaming and through capabilities used with the Hyper-V Extensible Switch.

Host Resiliency with NIC Teaming

NIC Teaming gives the ability to bond multiple high-speed network interfaces together into one logical NIC to support workload applications that require heavy network I/O and redundancy (Figure 20). Windows Server 2012 offers fault tolerance of network adapters with inbox NIC Teaming. This provides advanced networking capabilities to aggregate bandwidth from multiple network adapters and traffic failovers to prevent connectivity loss (so that failure of one NIC within the team does not affect the availability of the workload).

Figure 20: NIC Teaming in a virtual machine configuration

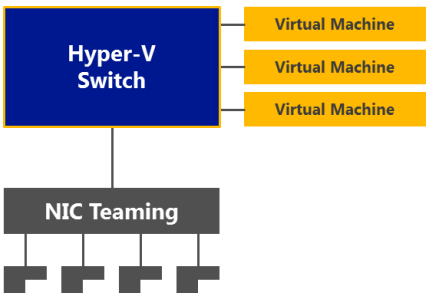

The built-in NIC Teaming solution in Windows Server 2012:

- Works with all network adapter vendors.
- Eliminates potential problems caused by proprietary solutions.
- Provides a common set of management tools for all adapter types.
- Is fully supported by Microsoft.

The solution also works within a virtual machine hosted on Hyper-V by allowing virtual network adapters to connect to more than one Hyper-V switch and still have connectivity even if the NIC underlying that switch gets disconnected. NIC Teaming uses two basic sets of configuration algorithms to provide better flexibility when designing networking for complex scenarios: switch-dependent mode and switch-independent mode.

Switch-Dependent Mode: These algorithms require all the network adapters of the team to be connected to the same switch. Two common ways in which the switch-dependent mode can be configured are as follows:

- Generic, or static, teaming (IEEE 802.3ad draft v1) requires configuration on the switch and computer to identify which links form the team.
- Dynamic teaming (IEEE 802.1ax, LACP) uses the Link Aggregation Control Protocol (LACP) to dynamically identify links between the computer and a specific switch.

Switch-Independent Mode: These algorithms do not require the switch to participate in the teaming. The team network adapters can be connected to different switches because a switch does not know to which network adapter it belongs.

Best Practices and Recommendations

We recommend that you use host-level NIC Teaming to increase resiliency and bandwidth. NIC Teaming supports up to 32 NICs from mixed vendors. It is important to have NICs within a team with the same speed.

Hyper-V Extensible Switch

Shown in Figure 21, the Hyper-V Extensible Switch is a layer-2 virtual interface that provides programmatically managed and extensible capabilities to connect virtual machines to the physical network.[23] With its new features and enhanced capabilities, the Hyper-V Extensible Switch supports tenant isolation, traffic shaping, protection against malicious SQL Server virtual machines, and simplified troubleshooting. With built-in support for Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers, the Hyper-V Extensible Switch also enables independent software vendors (ISVs) to create extensible plug-ins (known as *Virtual Switch Extensions*) that can provide enhanced networking and security capabilities.

Figure 21: Hyper-V Extensible Switch

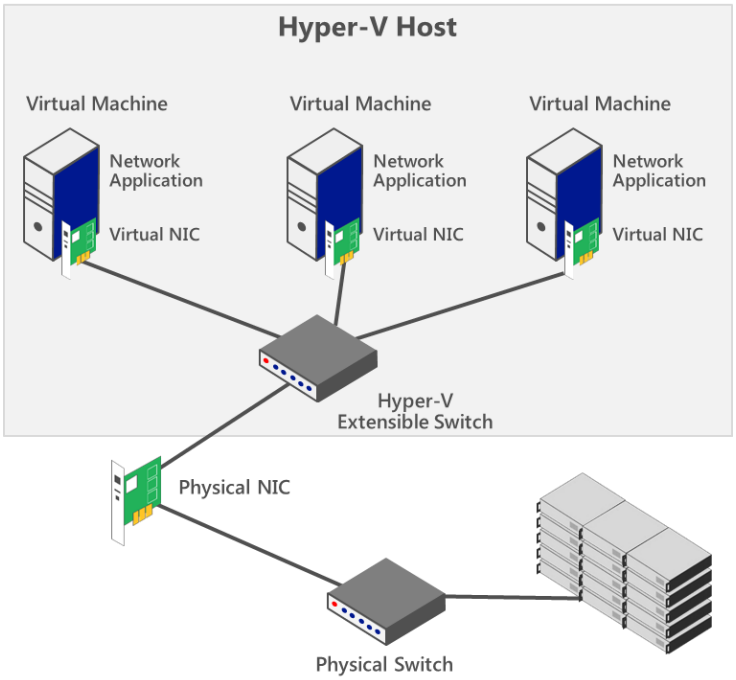

The two public Windows platforms for extending Windows networking functionality are used as follows:

- **NDIS Filter Drivers**: Used to monitor or modify network packets in Windows.
- **WFP Callout Drivers**: Used to allow ISVs to create drivers to filter and modify TCP/IP packets, monitor or authorize connections, filter IPsec-protected traffic, and filter RPCs. Filtering and modifying TCP/IP packets provides unprecedented access to the TCP/IP packet processing path. In this path, the outgoing and incoming packets can be modified or examined before additional processing occurs. By accessing the TCP/IP processing path at different layers, firewalls, antivirus software, diagnostic software, and other types of applications and services can be easily created.

Extensions can extend or replace the following three aspects of the switching process: ingress filtering, destination lookup and forwarding, and egress filtering.

Table 3 lists top companies that offer networking extensions for Hyper-V environments, as well as their key extension products.

Table 3: Options for networking extensions

CISCO	Cisco Nexus® 1000V Series Switches and Cisco Unified Computing System™ Virtual Machine Fabric Extender (VM-FEX)
NEC	NEC ProgrammableFlow PF1000
inMon The Inventors of sFlow®	InMon sFlow Agent
5NINE SOFTWARE	5nine Security Manager for Hyper-V

Virtual LANs

Virtual LANs (VLANs) subdivide a network into logical groups that share common physical infrastructure to provide network isolation. A VLAN uses explicit tagging in the Ethernet frames, and relies on Ethernet switches to enforce isolation and restrict traffic to network nodes of the same tag. However, there are some drawbacks with VLANs that limit networking capabilities within a large and complex network that provides communications for mission-critical workloads.

Windows Server 2012 introduces support for private VLANs (PVLANs) that extends the VLAN capabilities by providing isolation between two virtual machines on the same VLAN. Network virtualization in Windows Server 2012 removes the constraints of VLAN and hierarchical IP address assignment for virtual machine provisioning. Windows Server 2012 PVLANs provide scalability and better isolation of workloads. With PVLANs, a VLAN domain can be divided into subdomains that are represented by a pair of VLANs (primary VLAN and secondary VLAN). In such an implementation, every virtual machine in a PVLAN is assigned one primary VLAN ID and one or more secondary VLAN IDs. There are three modes for secondary PVLANs (Figure 22):

- **Isolated**: Isolated ports cannot exchange packets with each other at layer 2. If fact, isolated ports can only talk to promiscuous ports.
- **Community**: Community ports on the same VLAN ID can exchange packets with each other at layer 2. They can also talk to promiscuous ports. They cannot talk to isolated ports.
- **Promiscuous**: Promiscuous ports can exchange packets with any other port on the same primary VLAN ID (secondary VLAN ID makes no difference).

Figure 22: PVLAN in Windows Server 2012

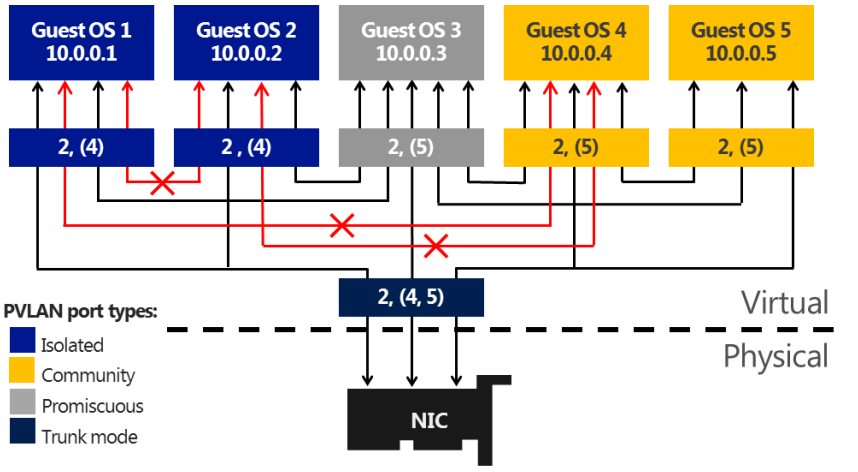

Example PVLAN:
– Primary VLAN ID is 2
– Secondary VLAN IDs are 4 and 5

Converged Networking through Data Center Bridging

Data Center Bridging (DCB) is based on a suite of standards from the Institute of Electrical and Electronics Engineers (IEEE) that enable converged fabrics in the data center, where storage, data networking, cluster IPC, and management traffic all share the same Ethernet network infrastructure. DCB in Windows Server

2012 provides networking enhancements by addressing congestion and reserving bandwidth in the NIC and the network for specified marked data types.

A DCB converged fabric reduces future TCO and simplifies the management of a virtualized environment. DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control. Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support iSCSI, RDMA over Converged Ethernet, or FCoE. DCB helps to provide lossless transport for mission-critical workloads by assigning priority-based flow control, Quality of Service (QoS), and guaranteed minimal bandwidth.

Figure 23 shows an example of a Windows Server 2012 cloud infrastructure with a highly available Hyper-V cluster that hosts virtual machines.[24] The virtual machines can be managed to create private or public clouds that are connected to a converged 10-GB Ethernet network and use dedicated file servers as the storage nodes. The converged networking infrastructure provides physical segmentation of infrastructure and tenant traffic. Each computer in the Hyper-V failover cluster must have at least two network adapters so that one adapter can host the cloud infrastructure traffic and one adapter can support tenant traffic. To make the NICs more resilient against NIC failures, additional network adapters can be added on each of the networks and teamed using Windows Server 2012 Load Balancing and Failover NIC Teaming. The NICs can be 10 GbE or 1 GbE network adapters. These NICs are used for live migration, cluster, storage, and management (together referred to as *infrastructure traffic*) as well as tenant traffic. By creating separate NIC teams, you can physically separate cloud infrastructure traffic from cloud tenant traffic.

Figure 23: High-level view of Windows Server 2012 cloud infrastructure

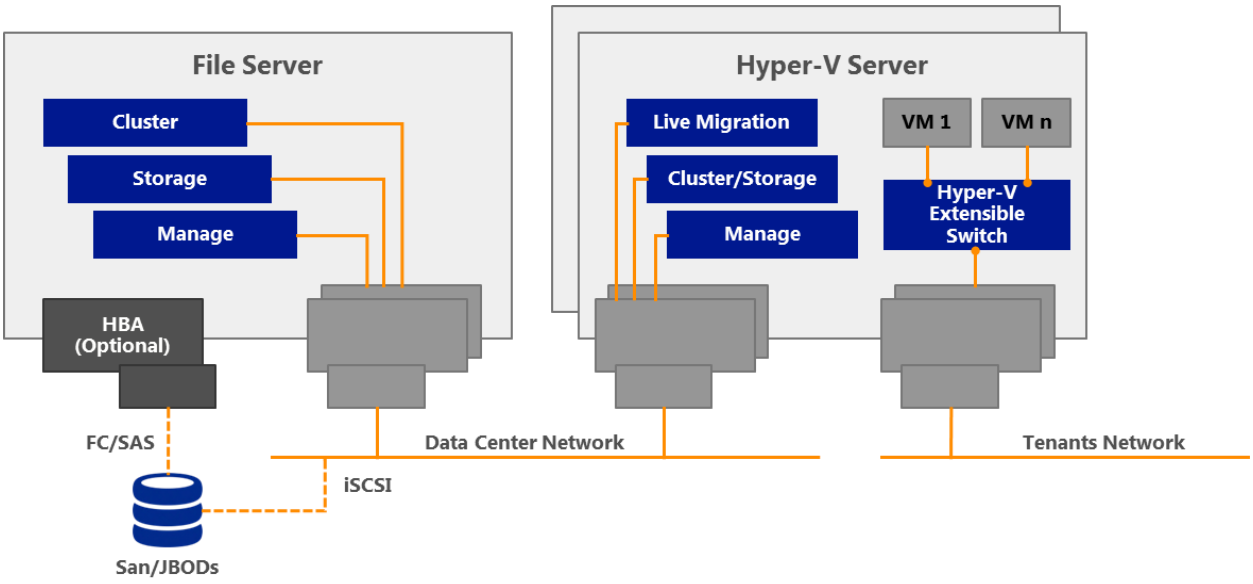

Windows QoS policies on the parent Hyper-V partition enable guaranteed levels of bandwidth to all incoming infrastructure traffic. DCB enables QoS offload capabilities to both the 10 GbE network adapter and the 10 GbE-capable switch. The tenant traffic from different tenants receives guaranteed levels of bandwidth using Hyper-V virtual switch QoS policies.

Host Resiliency and Virtual Machine Agility

High availability of databases is becoming more important as organizations strive to make mission-critical data and applications available to end users. Windows Server 2012 provides capabilities to enhance the availability and scalability of SQL Server 2012 instances. Further, Windows Server 2012 virtualization provides a highly available and scalable platform for SQL Server 2012.

Host Clustering

This subsection discusses the key elements of host clustering, including failover clustering, Cluster Shared Volumes, clustering recommendations, virtual machine priority, virtual machine affinity, and live migration.

Failover Clustering

Failover clustering allows you to connect physical machines (also called nodes) together to provide better scalability and high availability. These clustered nodes work in such a way that if one or more of the active nodes fail, the other nodes in the cluster begin to provide service (Figure 24). Clustered nodes are continuously monitored to ensure that they are working properly. The nodes come to know each other's active status by using a heartbeat—a periodic signal between two directly connected machines.

Figure 24: Failover clustering—virtual machines fail over to Node 1 simultaneously

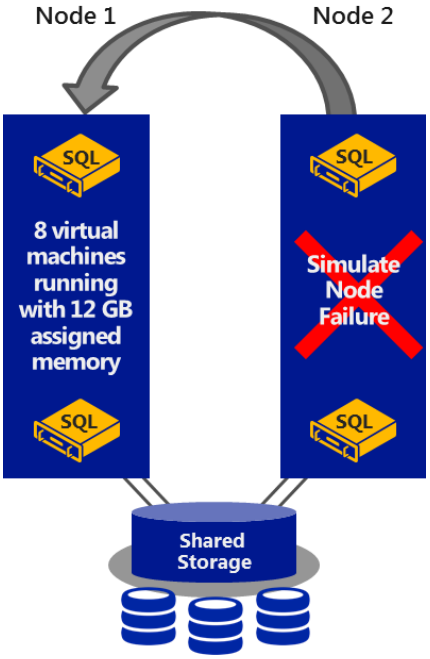

Best Practices and Recommendations

We recommend that all hardware is certified for Windows Server 2012 and that the complete failover cluster solution passes all tests in the Validate a Configuration Wizard. For more information about validating a failover cluster, see Validate Hardware for a Windows Server 2012 Failover Cluster.

Windows Server 2012 Hyper-V supports scaling clusters up to 64 nodes and 8,000 virtual machines per cluster. Windows Server 2012 also provides Windows PowerShell cmdlets and a snap-in for Failover Cluster Manager, which allows administrators to manage multiple clustered nodes.

Cluster-Aware Updating

In the past, it has been a challenging task for administrators to appropriately update and patch failover clusters. The Cluster-Aware Updating (CAU) feature in Windows Server 2012 simplifies this task. CAU facilitates automated maintenance of cluster nodes/servers. Automating cluster nodes makes the server maintenance process in a cluster faster, easier, more reliable, and more consistent with less downtime. CAU puts each cluster node in maintenance mode, applies required updates/patches, and restores the node to be used again. At a high level, CAU performs the following steps:[25]

- Puts a node of the cluster in maintenance mode and takes it offline transparently.
- Moves clustered roles off the node.
- Installs the updates or patches, and any dependent updates.
- Performs a restart, if needed.
- Brings the node back online and out of maintenance mode.
- Restores clustered roles on the node.
- Moves to the next node and updates/patches it in the same manner.

This increases the availability of servers during the update and patching process in both environments (virtualized and non-virtualized). It also helps to maintain the security and performance of servers in the cluster. Administrators use the Cluster-Aware Updating Wizard for automating the update of a failover cluster (Figure 25).

Figure 25: Cluster-Aware Updating Wizard

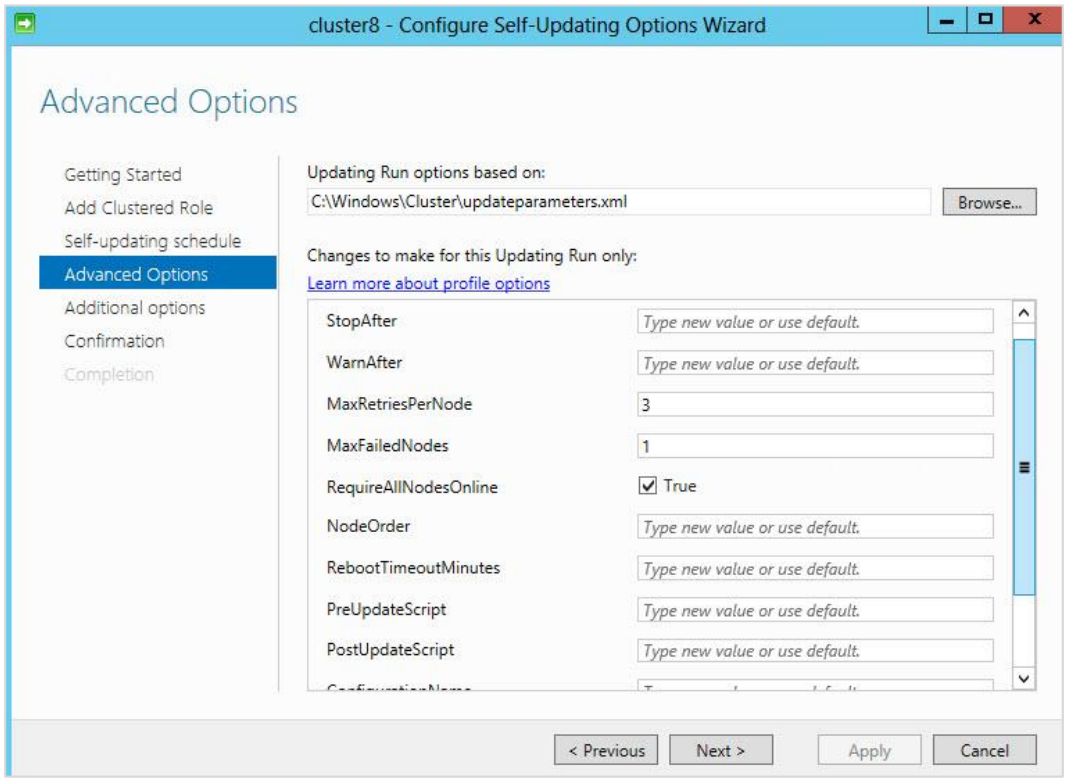

CAU can perform the cluster updating process in two different modes: self-updating mode and remote-updating mode. In self-updating mode, the CAU clustered role is configured as a workload on the failover cluster that is to be updated. In remote-updating mode, a remote computer running Windows Server 2012 or Windows 8 is configured with the CAU clustered role. This remote computer is also called the *Update Coordinator* and is not part of the cluster that is updating.

Note that for many clustered roles in a cluster, the automatic updating process triggers a planned failover, which, in turn, can cause a service interruption for a very short time (transient).[26]

Best Practices and Recommendations

Use the CAU feature with continuously available cluster workloads in Windows Server 2012 to perform updates on the clusters with no impact on service availability. Examples of continuously available cluster workloads in Windows Server 2012 are file servers (file server with SMB Transparent Failover) and Hyper-V with live migration.

Create *Updating Run Profiles* for different classes of failover clusters, and store and manage them on a centralized file share. This ensures that the CAU deployments consistently apply updates to the clusters throughout the IT organization (even across different departments, line-of-business areas, or administrators).

CAU supports an extensible architecture that helps to update the cluster node with node-updating tools and software updates that are not available from Microsoft or through Windows Update or Microsoft Update. Examples include custom software installers, updates for non-Microsoft device drivers, and network adapter/HBA firmware updating tools. This is beneficial for publishers who want to coordinate the installation of non-Microsoft software updates.

Cluster Shared Volumes

With Cluster Shared Volumes (CSV), Windows Server 2012 provides shared storage resources for clustered virtual machines. The CSV technology makes it easy to manage large numbers of logical unit numbers (LUNs) in a failover cluster. CSV allows simultaneous read-write permissions to the same NTFS-provisioned LUN disk (Figure 26). Due to significant improvements in CSV design in Windows Server 2012, Cluster Shared Volumes can perform an increased number of operations in direct I/O mode, as compared to previous versions of Windows Server. This feature can also be helpful while operating workloads for failover, conducting live migration, or moving workloads; it allows workloads to share the same volume without impacting the other workloads.

- CSV provides better security for mission-critical data by supporting BitLocker Drive Encryption.
- For decrypting data, CSV uses cluster name objects as identity for data decryption.
- CSV provides continuous availability and scalability of file-based server storage for applications or databases using the Scale-Out File Server feature.
- CSV supports two types of snapshots, including application-consistent and crash-consistent Volume Shadow Copy Service snapshots.
- CSV supports clustered VHD files for clustered workloads.

Figure 26: Cluster Shared Volumes

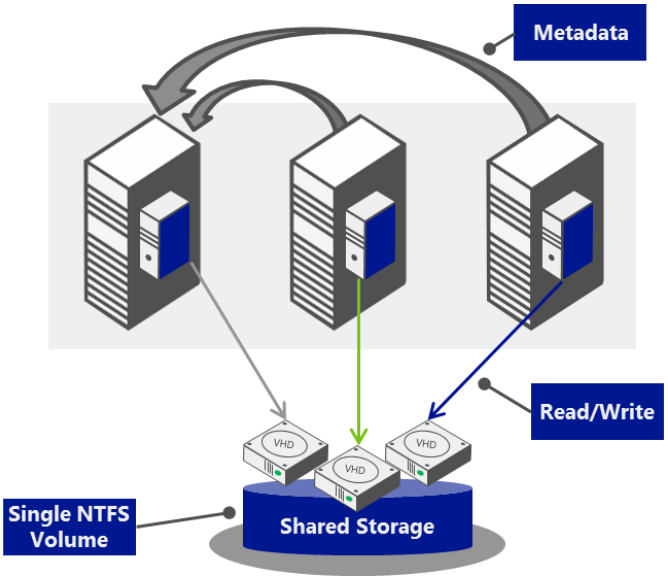

Best Practices and Recommendations

The following guidelines are recommended for deploying a SAN with a failover cluster:

- Confirm with the manufacturers and vendors whether the storage (including drivers, firmware, and software) being used is compatible with failover clusters in Windows Server 2012.
- Isolate the LUN used for one set of cluster servers from all other servers through LUN masking or zoning.
- Consider multipath I/O software for the highest level of redundancy and availability.

Clustering Recommendations

This subsection highlights best practices and recommendations around networking, storage, and compute when designing and configuring the cluster networks in a data center/private cloud.[27]

Networking

The following are best practices and recommendations for networking:

Best Practices and Recommendations

If you want your existing network architecture (built on Windows Server 2008 R2 Hyper-V) to be the same in terms of network traffic segmentation and isolation in Windows Server 2012 Hyper-V: Install separate physical NICs for each type of traffic (such as infrastructure traffic and tenant traffic). Also, assign VLAN 802.1q tags to each adapter.

To dedicate each traffic type to a specific adapter, configure each of the traffic flows to use the correct subnet/IP on the dedicated NIC.

To achieve the highest network performance possible, use Single Root I/O Virtualization (SR-IOV). This is a new feature in Windows Server 2012 that enables virtual machines to directly access network adapter hardware and bypass the virtual networking stack.

In a non-converged infrastructure that supports multiple traffic types, have enough network adapters on each system in the failover cluster to support each traffic type. This is especially important when you need to isolate traffic types from each other. It includes assigning network adapters for infrastructure traffic (live migration, cluster/CSV, and management) and storage traffic. Based on your requirements, infrastructure traffic network adapters can be non-teamed 1 GB adapters, 10 GB adapters, or a mix of 1 GB and 10 GB adapters. However, a tenant should have a 10 GB adapter supporting SR-IOV.

Storage

The following are best practices and recommendations for storage:

Best Practices and Recommendations

If your organization already has FC or iSCSI SAN and requires all virtual machines in a failover cluster to be connected to the block storage, you can do the following:

- Configure Hyper-V hosts to keep using traditional SAN storage.
- Enable each member of the Hyper-V failover cluster to connect to iSCSI, FC, or FCoE to access the block storage.

Compute

The following are best practices and recommendations for compute:

Best Practices and Recommendations

To use the same hardware that was used for running Hyper-V on Windows Server 2008 R2, add a 10 GbE network adapter to support SR-IOV. Or use a 1 GbE adapter and do not deploy SR-IOV.

To ensure the highest virtual machine density possible per host server, use processor offload technologies such as Receive-Side Scaling (RSS), DCB, Receive-Side Coalescing (RSC), and RDMA.

DCB and RDMA cannot be used where there is no dedicated storage NIC. This is because these features require direct access to hardware.

Virtual Machine Priority

IT administrators can configure availability options for virtual machines running on Hyper-V host clusters. An administrator sets priority for the virtual machines in a host cluster, which the host cluster uses to identify the high-priority virtual machines and give them first preference. This ensures that high-priority virtual machines are allocated memory and other resources first for better performance.[28]

In Windows Server 2012, administrators can configure availability options/settings to provide improved and efficient allocation of cluster resources (such as when starting or maintaining nodes) in large physical clusters and Hyper-V failover clusters. The availability options for managing clustered virtual machines and other clustered roles include:[29]

Priority Settings: This option can be applied to all clustered roles, including clustered virtual machines. A virtual machine can be set to high priority, medium priority, low priority, or No Auto Start. By default, every virtual machine is set to medium priority. Clustered roles with higher priority are started and placed on nodes before those with lower priority. If a No Auto Start priority is assigned, the role does not come online automatically after it fails, which keeps resources available so other roles can start.

Preemption of Virtual Machines Based on Priority: This option can be applied to clustered virtual machines. In case of a node failure, if the high-priority virtual machines do not have the necessary memory and other resources to start, the lower priority virtual machines are taken offline to free up resources. When necessary, preemption starts with the lowest priority virtual machines and continues to higher priority virtual machines. Virtual machines that are preempted are later restarted in priority order.

Best Practices and Recommendations

SQL Server virtual machines are typically the backend for other applications or services. Therefore, these virtual machines usually need to start first.

Virtual Machine Affinity

The Virtual Machine Affinity rules in Windows Server 2012 allow administrators to configure partnered virtual machines to migrate simultaneously at failover. For example, let's say there are two machines that are partnered: One virtual machine has front-end applications and the other has a back-end database. These two virtual machines can be configured to always fail over together to the same node.

It can also be specified that two particular virtual machines cannot coexist on the same node in a failover scenario. This is the anti-affinity virtual machine rule. In this case, Windows Server 2012 Hyper-V migrates the partnered virtual machines to different nodes to help mitigate a failure. For example, domain controllers running on the same Hyper-V host can be migrated to different nodes to prevent loss of the domain in case of failure. Windows Server 2012 Hyper-V provides a cluster group property called *AntiAffinityClassNames* that can be applied to any virtual machine in the Hyper-V cluster group. This property allows preferences to be set to keep a virtual machine off the same node as other virtual machines of a similar kind.

Best Practices and Recommendations

Use Virtual Machine Affinity rules to keep SQL Server virtual machines apart on hosts.

Live Migration

Live migration is a process where running workloads can be moved from a source server to a destination server without impacting the availability of running business applications or critical data. While migrating live virtual machines, there are two major concerns: outage of applications or data, and prevention of data loss.

Windows Server 2012 with Hyper-V provides a better way to migrate running virtual machines from one physical server to another without hampering business availability. Hyper-V 2012 with the enhanced Live Migration feature allows you to execute the live migration of multiple workloads at the same time. For SQL Server 2012, Hyper-V offers a simplified way to migrate clustered as well as non-clustered SQL Server workloads without any outage. During the migration process of any workload, no additional configuration changes are required on the guest operating system.

- Windows Server 2012 with Hyper-V allows the use of up to 10 GB network bandwidths to make migration faster.
- Hyper-V offers functionality to migrate multiple virtual SQL Servers simultaneously.

For migrating SQL Server 2012, Hyper-V sets up a TCP connection between source and destination SQL Server virtual machines to transfer virtual machine configurations. The memory assigned to the migrating SQL Server virtual machine is transferred over the network to the destination SQL Server virtual machine. Hyper-V keeps track of the memory pages being modified during the transfer from the source to the destination server. Once all the modified pages are copied completely to the destination server, Hyper-V takes care of the migration of any associated virtual hard disk files or physical storage attached through a virtual FC adapter. After completing all stages of migration, Hyper-V brings up the new destination SQL Server virtual machine (Figure 27).

Figure 27: Live migration with Hyper-V

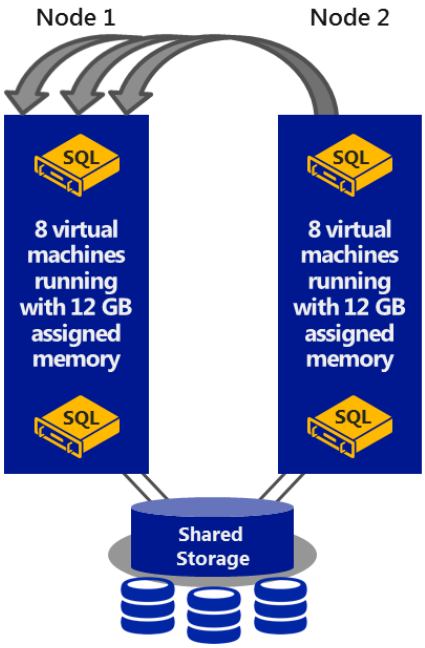

In addition to improvements made with the Live Migration feature, Windows Server 2012 now allows the live migration of virtual machine storage—independent of the virtual machine itself and without any downtime. This is known as *live storage migration* and can be initiated using the Hyper-V Manager console, Failover Cluster console, Microsoft System Center Virtual Machine Manager (SCVMM) console, or PowerShell. This capability can be used in a number of scenarios, including a need to redistribute virtual machine storage when disk space may run low or in order to perform maintenance on underlying storage without disrupting the virtual machine. Storage live migration for .vhd and .vhdx disks can be performed on client Hyper-V, standalone Hyper-V server, and clustered Hyper-V servers. Storage live migrations of virtual machines do not support local pass-through disks.

Windows Server 2012 also implements *shared-nothing live migration* that helps in migrating virtual machines using just a network cable without any downtime. Shared-nothing live migration works in conjunction with live migration and live storage migration. First the storage is live migrated. Once the live storage migration completes, the state of the virtual machines is copied and synchronized between the source and the destination servers over the network.

Hyper-V Replica

For maintaining high availability of mission-critical applications, organizations can replicate their applications on different server workloads in different locations. Hyper-V with Windows Server 2012 provides Hyper-V Replica for replicating virtual machines from one location to another location. Hyper-V Replica is not limited to any specific type of hardware vendor for site replication. It provides software-based, asynchronous replication of virtual machines between sites. Hyper-V Replica can help to replicate running workloads in a minimal span of time by ensuring maximum availability of SQL Servers, databases, agents, and other SQL Server components.

- Hyper-V Replica can be implemented in a standalone server model, failover cluster model, or mixed mode environment of both.
- Hyper-V Replica allows you to store VHD files in different locations to make them accessible in case a data center goes down for any reason.

Hyper-V Replica works with any IP-based network to replicate workloads. It provides a secure data transmission through optional encryption. Hyper-V Replica uses two sites: primary and secondary. The primary site is where applications such as SQL Server are currently running, and the secondary site is the destination site where the workloads need to be replicated. If the primary site goes down, the secondary site, which replicated asynchronously, can act as the primary site for high availability (Figure 28).

Figure 28: Hyper-V Replica, primary and secondary sites

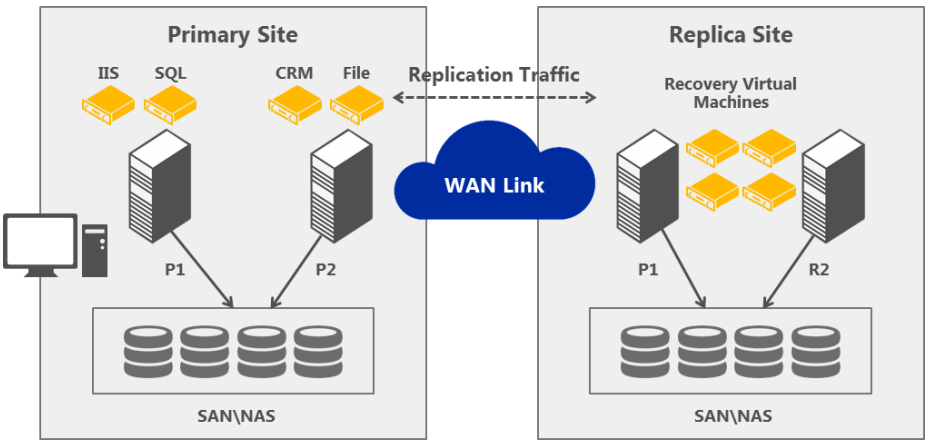

Best Practices and Recommendations

Key SQL Server considerations on Hyper-V Replica include:[30]

- SQL Server on Hyper-V Replica is supported as long as the flag for *EnableWriteOrderPreservationAcrossDisks* is set.
- If multiple SQL Server virtual machines are tightly coupled with one another, individual virtual machines can fail over to the disaster recovery site, but SQL Server high availability features need to be removed and reconfigured. For this reason, the following SQL Server features are not supported on Hyper-V Replica: availability groups, database mirroring, failover cluster instances, log shipping, and replication.

Virtual Machine Configuration

In addition to configuring and establishing the host server as a virtualization server with Hyper-V, it is important to design detailed architecture and system specifications for building virtual machines for expected workloads. It is also necessary to plan for needed resources for the virtual machines. The number of virtual machines you can run on any individual server depends on the server's hardware configuration and the anticipated workloads.

Microsoft Assessment and Planning Toolkit

Because databases are the platforms for a wide variety of applications, they tend to spread across organizations, leading to a proliferation of database instances also known as *database sprawl*. Therefore, databases are one of the prime candidates for consolidation.[31] Consolidation of databases reduces the needed space and capital and operating expenditures in a data center. It helps to improve efficiency and agility for better control and flexibility of computing resources in terms of their placement, sizing, and overall utilization. Consolidation enables standardization of requirements and methodologies and centralized control for sharing common schemas and management infrastructure. Consolidation also helps to reduce power and thermal costs, thereby providing a "greener" operating environment.

The Microsoft Assessment and Planning (MAP) toolkit provides a complete network-wide inventory of SQL Server 2012 for comprehensive process planning and for migrating legacy database instances. MAP provides information that can be used to consolidate databases and better utilize hardware and database resources. Besides this, MAP helps in gathering information on resource usage for SQL Server consolidation, captures IOPS, and provides detailed reporting for the requirements of database deployments in a virtualized environment.

SQL CPU Considerations

This subsection discusses Non-Uniform Memory Access from the virtual machine perspective. Virtual NUMA is the foundation for several key CPU considerations for SQL Server.

Non-Uniform Memory Access – Virtual Machine Perspective

In Windows Server 2012, NUMA is now extended to the virtual environment (virtual NUMA) by making virtual NUMA topology available to the guest operating systems. High-performance applications such as SQL Server support NUMA and use the computer's NUMA topology to increase performance by considering NUMA when scheduling threads or allocating memory. Therefore, reflecting the underlying NUMA topology into virtual machines running SQL Server 2012 reduces the need for remote memory access, which is critical to improving SQL Server 2012 performance.

To identify and adapt to the virtual NUMA topology within the virtual machines, the NUMA-aware guest operating system and applications use their inherent NUMA performance optimizations. In this way with Windows Server 2012 Hyper-V, the default virtual NUMA topology is optimized to match the NUMA topology of the host/physical computer, as shown in Figure 29.[32]

Figure 29: Guest NUMA topology by default matching host NUMA topology

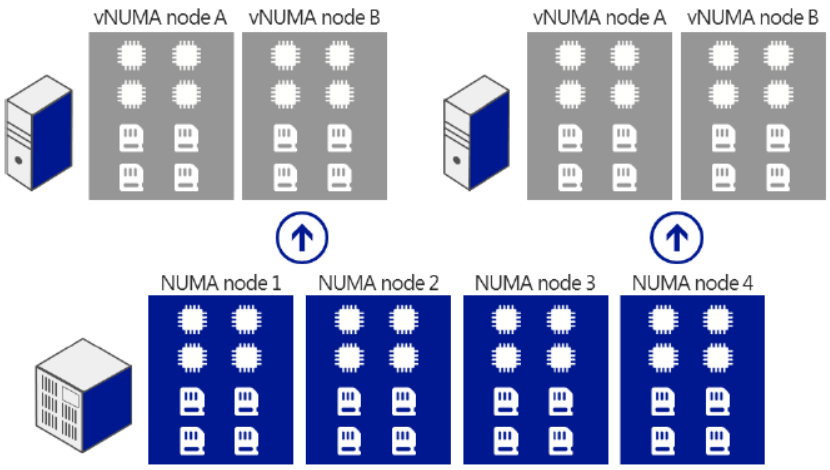

The best practices below provide more guidance around managing varying CPU demand, reducing overhead on the CPU, and optimizing processor performance for SQL Server workloads.[33, 34]

Best Practices and Recommendations

Hyper-V publishes performance counters like Performance Monitor (Perfmon.exe) and Logman.exe. These performance counters help to characterize the behavior of the virtualization server and report resource usage. To measure CPU usage of the physical host, use the Hyper-V Hypervisor Logical Processor performance counters. The Performance Tuning Guidelines for Windows Server 2012 contain the list of performance counters to monitor performance.

Virtual machines with multiple virtual processors have additional overhead related to synchronization costs in guest operating systems. Therefore, if a virtual machine will never have CPU-intensive loads, even at peak hours, configure it to use only one virtual processor. Multiple virtual processors should only be configured in cases where the virtual machine requires more processing power under peak loads.

Identify and categorize virtual machines based on the intensity of the loads they bear (high-intensity loads and low-intensity loads). Then set weights and reserves on the virtual processors accordingly. In this way, you can ensure that a large amount of the CPU cycle is available for virtual machines/virtual processors having high-intensity loads.

Whenever possible, avoid the use of emulated devices for SQL Server deployments as this can lead to significant CPU overhead. Install the latest virtual machine Integration Services in each supported guest virtual machine. Virtual machine Integration Services helps to improve I/O throughput and decrease overall CPU usage of guests. This is because it includes enlightened drivers for Hyper-V-specific I/O devices that reduce CPU overhead for I/O.

SQL Memory Considerations

For memory in a virtualized environment, better performance and enhanced support are essential considerations. You must be able to both quickly allocate memory to virtual machines depending on their requirements (peak and off-peak loads) and ensure that the memory is not wasted. New enhancements in Windows Server 2012 help to optimize the utilization of memory allocated to virtual machines.[35]

Dynamic Memory

Earlier versions of Hyper-V allowed administrators to assign a fixed amount of physical memory to a virtual machine on the host machine. Once the memory was assigned, it was not possible to change the memory for that particular virtual machine during its run state.[36, 37] To overcome this problem, Microsoft introduced the concept of Dynamic Memory in Windows Server 2008 R2 SP1.

With Windows Server 2012, Microsoft has enhanced the Dynamic Memory feature to provide increased agility in how the memory is allocated and managed between virtual machines running on a host. Dynamic Memory in Windows Server 2012 has introduced two key new enhancements: minimum memory and Hyper-V smart paging.

Minimum Memory

Minimum memory allows Hyper-V in Windows Server 2012 to reclaim the unused memory from virtual machines. This results in increased virtual machine consolidation. However, there can be a limitation to this feature. If you must restart one virtual machine and it has less memory than required for its startup memory, Hyper-V needs additional memory to restart the machine. Yet, Hyper-V may not always have additional memory available. Such a situation results in a virtual machine start failure. To overcome this situation, Dynamic Memory in Windows Server 2012 has introduced Hyper-V Smart Paging.

Hyper-V Smart Paging

Hyper-V Smart Paging is a memory management technique that is used to cover the gap between minimum memory and startup memory, enabling reliable restart of virtual machines (Figure 30). It uses disk resources as additional, temporary memory when more physical memory is required to restart a virtual machine than is currently available.[38, 39]

Figure 30: Hyper-V Smart Paging

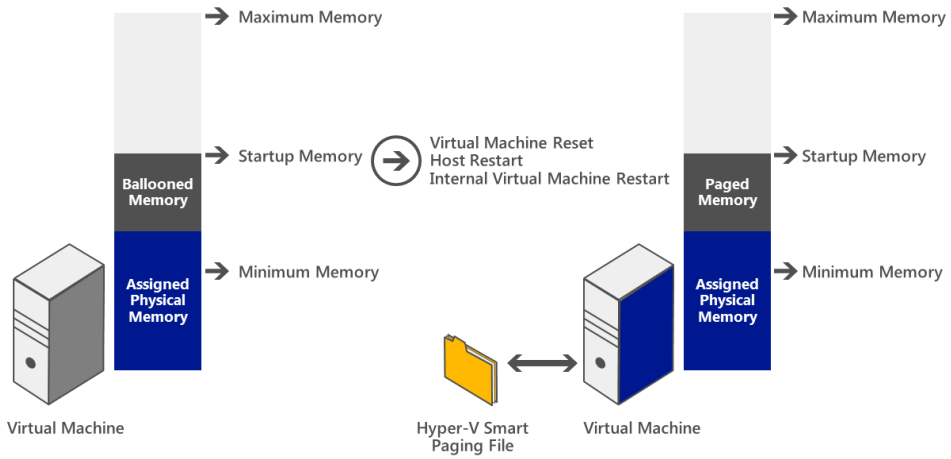

Hyper-V Smart Paging can lead to some performance degradation due to slower disk access speeds. Therefore, to ensure that the performance impact of Smart Paging is minimized, this feature is used only when all of the following are true:

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

Best Practices and Recommendations

In environments where performance is critical, use SSD for Smart Paging.

Memory Ballooning: Memory ballooning is a technique used to further reduce the performance impact of Hyper-V Smart Paging. Once the virtual machine is restarted and the need for memory is less than the startup memory, Hyper-V can stop using Smart Paging. Therefore, Hyper-V removes the temporary memory from the virtual machine by coordinating with Dynamic Memory components inside the guest (a process sometimes referred to as *ballooning*). With this technique, use of Hyper-V Smart Paging is temporary and is not expected to be longer than 10 minutes. Figure 31 shows Hyper-V removing memory from the virtual machine after it completes the startup process.[40]

Figure 31: Removing paged memory after virtual machine restart

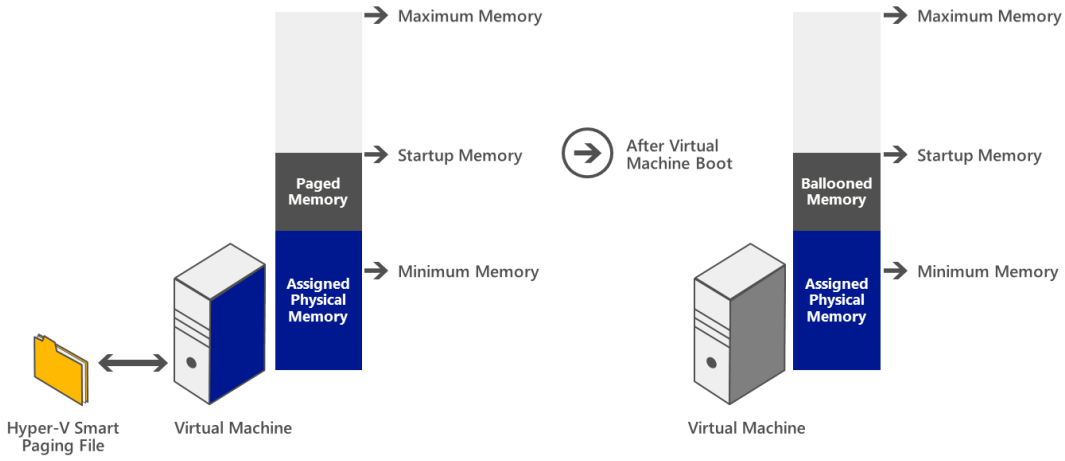

The best practices below provide more guidance around planning and managing memory for virtual machines running SQL Server 2012 workloads.[41, 42, 43]

Best Practices and Recommendations

Allocate a reasonable amount of memory to the virtual machines running SQL Server workloads so that they can handle the expected loads (at peak and off-peak times). If the memory is not sufficient, it can increase response time or I/O usage for highly intensive workloads. Allocate at least 512 MB of memory for the virtual machine running Windows Server 2012.

Best Practices and Recommendations

Be sure to check the minimum memory requirement of the SQL Server workload that will be hosted on a Windows Server 2012 guest machine and, based on that, allocate the total minimum memory (which is likely to be much higher than 512 MB). Dynamic Memory can be enabled to manage the memory requirements dynamically.

There is a limit to the amount of physical memory with which a physical server can be equipped. Therefore, prioritize memory usage based on the requirements of the running virtual machines or the importance of SQL Server 2012 workloads, such as OLTP, DW, and BI. Assign memory weights to high-priority virtual machines/workloads so that in the event of a physical memory crunch, memory is first allocated to the higher priority virtual machines.

In general, it is better to leave max server memory at its default setting. This allows SQL Server to manage memory dynamically.

To provide better stability to a virtual machine workload, grant *Lock Pages in Memory* user rights to the SQL Server service account. This helps when Hyper-V Dynamic Memory is trying to reduce the virtual machine's memory. In such cases, it prevents Windows from paging out a large amount of buffer pool memory from the process, thereby providing a positive performance impact.

If memory pre-allocation is used with SQL Server, it is better to assign only the amount of memory that is required so sufficient memory remains for other processes (thereby avoiding paging). This also ensures that enough memory is available for Analysis Services. (Use the peak value for the Process: Private Bytes counter for the msmdsrv instance to establish this value).

Do not set the memory pre-allocation value too high. Otherwise, other processes may not get sufficient memory at the time when they require it, and this can result in memory paging. If the SQL Server relational engine and Analysis Services are running on the same server, also limit the maximum memory that can be used by the SQL Server relational engine. It is generally sufficient to leave 6-8 GB of memory for the operating system if no other services are running on the same server. However, this can vary depending on workload requirements.

By default, a virtual machine gets its preferred NUMA node every time it runs. In due course, an imbalance in the assignment of NUMA nodes to the virtual machines may occur. This may happen because of the ad hoc memory requirements of each virtual machine and because the virtual machines can be started in any order. Therefore, we recommend that you use Perfmon to check the NUMA node preference settings for each running virtual machine. This can be checked through the *\Hyper-V VM Vid Partition (*)\ NumaNodeIndex* counter.

To automatically change NUMA node assignments depending on the requirements of running virtual machines, execute NUMA node balancing.

Crossing the NUMA boundary can reduce virtual performance by as much as 8 percent. Therefore, configure a virtual machine to use resources from a single NUMA node.

Best Practices and Recommendations

Reduce the max server memory prior to live migration. Use sp_configure or the Memory Weight option in Hyper-V Dynamic Memory to do so. Note that you can use sp_configure only if you have appropriate rights on SQL Server 2012.

Adjusting max server memory prior to migration provides following benefits:

- Highest success rate for migrations.
- Most predictable impact to SQL Server workload.
- Most even rebalancing of physical memory assigned to the virtual machines.

SQL Storage Considerations

This subsection discusses two key storage considerations for SQL Server: virtual disks and guest storage.

Virtual Disks

When considering virtual disks, it is important to know the capabilities and limitations of fixed-size VHDX, pass-through disks, and virtual IDE/virtual SCSI. These three topics are discussed below.

Fixed-Size VHDX

Fixed-size VHDX uses the full amount of space specified during virtual hard disk creation. However, the size of a fixed-size VHDX can be increased by using Hyper-V Manager or running a PowerShell script. Note that reducing the size is not supported. Fixed-size VHDX can deliver better throughput than dynamic VHDX. Dynamic VHDX can also be used; however, these disks are slower and tend to be larger than fixed-size VHDX. When dynamically expanding disks, you run the risk of storage oversubscription, which could lead to fragmentation. If you choose fixed-size VHDX for SQL storage, 2 TB is the maximum recommended VHD size. Fixed VHDXs are recommended for portability and snap-shooting.

Best Practices and Recommendations

Why use VHDX? Hyper-V in Windows Server 2012 introduces VHDX, a new version of the VHD format that is designed to handle current and future workloads. VHDX has a much larger storage capacity than the older VHD format. It also provides protection from data corruption during power failures and optimizes structural alignments to prevent performance degradation on new, large-sector physical disks. The VHDX format provides the following advanced capabilities that accommodate large workloads like databases:

- VHD storage capacity: Supports up to 64 TB.
- Provides additional data protection against power failures by logging updates to the VHDX metadata structures.

Best Practices and Recommendations

- Enables improved alignment of the VHD format to work well on large sector physical disks for storing databases.
- Stores custom metadata about the file to record operating system version or updates.
- Provides larger block sizes for dynamic and differencing disks to address the needs of the SQL Server databases.
- 4 KB logical sector virtual disk: Allows for increased performance when used by SQL Server applications designed for 4 KB sectors.[44]
- Allows physical storage devices to reclaim unused space by trimming into smaller file sizes. (Trim requires physical disks directly attached to a virtual machine or SCSI disks, as well as trim-compatible hardware.)

Pass-Through Disks

Pass-through disks are physical disks or LUNs that are attached to a virtual machine. They provide improved performance at a small cost. By configuring disks as pass-through disks for SQL Server in a Hyper-V virtual machine, you gain the highest performance but negate the flexibility of using VHDX files.[45] Fixed-size VHDX files provide near-parity in terms of performance, and they usually should be used instead of the less flexible pass-through disks.

Virtual IDE vs. Virtual SCSI

Virtual machines can be configured to use virtual IDE device controllers or virtual SCSI device controllers to connect virtual storage. When starting a virtual machine, the virtual IDE controller is used because the virtual SCSI disks require the operating system to be enlightened during boot-up. IDE is limited to 3 connected disks. (One port is retained for the DVD drive, which is required for updating the integration components.) Virtual SCSI, on the other hand, can have 64 connected disks per controller and 4 controllers per virtual machine, giving a total of 256 virtual SCSI disks per virtual machine. Virtual SCSI also supports hot-add/removal of disks, whereas virtual IDE disks do not.

Best Practices and Recommendations

Attach the SQL drives (or any secondary drive) to the Virtual SCSI controller for more flexibility.

Guest Storage

When considering guest storage, it is important to know the capabilities and limitations of SQL Server over SMB, guest iSCSI, and Hyper-V Virtual Fibre Channel. These three topics are discussed below.

SQL Server over SMB

Windows Server 2012 includes detailed SMB client performance counters that provide information about I/O size, I/O latency, IOPS, and so on. This lets a SQL Server database administrator or Hyper-V administrator analyze the performance of the SMB file shares where their data is stored. SQL Server 2012 databases can be stored using remote shared folders that use SMB Direct and SMB Multichannel on industry-standard network adapters. The SMB client and SMB server are both optimized for small random read-write I/O to increase performance for common server application workloads, such as OLTP. SMB also uses a large Maximum Transmission Unit (MTU) feature, enabled by default, to significantly improve performance in large sequential transfers, such as those for DW, in addition to database backup and restore operations.

SQL Server over SMB is flexible enough to support a variety of possible configurations, including Single-Node File Server, Dual-Node File Server, and Multi-Node File Server modes (Figure 32).

Figure 32: Flexible configuration options for SQL Server over SMB (1 = Single-Node, 2 = Dual-Node, 3 = Multi-Node)

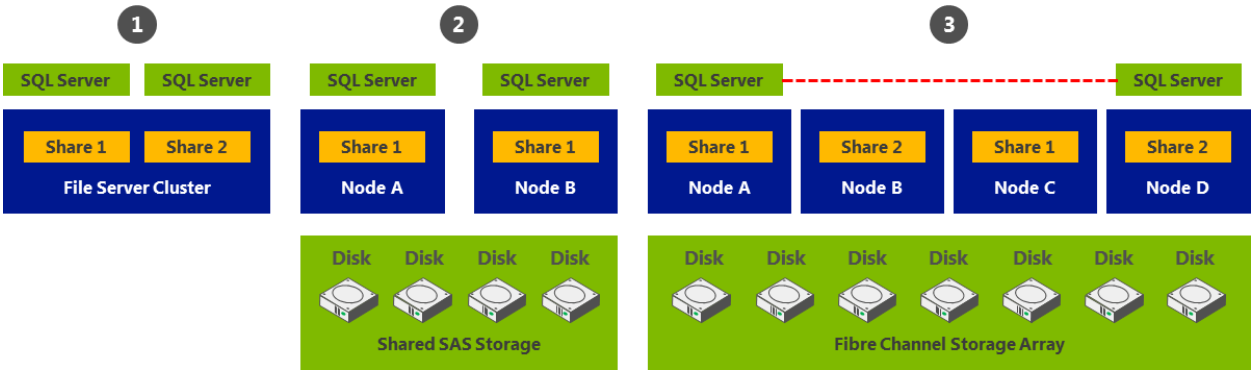

Table 4 compares the cost, availability, and scalability of the three configurations for SQL Server over SMB.

Table 4: Comparison of SQL Server over SMB configurations

	Single-Node File Server	Dual-Node File Server	Multi-Node File Server
Cost	Low cost	Medium cost	Higher cost
Availability	Shares not continuously available	Shares continuously available	Shares not continuously available
Scalability	Limited scalability (~100 spindles)	Medium scalability (~200 spindles)	Highest scalability (~1,000 spindles)

With the ability to take advantage of innovative SMB features like transparent failover, SMB Multichannel, and SMB Direct, organizations can now deploy storage for SQL Server 2012 workloads on high-performance Windows Server 2012 file servers that are cost efficient and continuously available.

Guest iSCSI

Deploying a SQL Server 2012 database on iSCSI storage provides a cost-effective solution for enterprise-level database deployments. The iSCSI target provides storage to SQL Server; SQL Server acts as the iSCSI initiator that consumes storage. SQL Server can store its program files, logs, and other data on iSCSI disk volumes in both cluster and non-cluster configurations.

Virtual Fiber Channel for Hyper-V

Virtual Fiber Channel for Hyper-V helps to connect to FC storage from within a virtual machine. It provides direct SAN access from the guest operating system by using a standard World Wide Name (WWN) associated with a virtual machine. Virtual FC for Hyper-V also helps to run the Failover Clustering feature inside the guest operating system of a virtual machine that is connected to the underlying, shared FC storage.

Best Practices and Recommendations

Use MPIO inside the guest with Virtual FC to ensure resilient connections from the virtual machine to storage.

For virtualizing SQL Server 2012, Virtual FC for Hyper-V allows you to use existing FC investments and provides support for live migration and MPIO.

Support for Live Migration: This support allows for the live migration of SQL Server virtual machines or instances across Hyper-V hosts while maintaining FC connectivity. It ensures that all LUNs are available on the destination host before migration, thereby helping to eliminate downtime.

MPIO Connectivity: The MPIO functionality with Virtual FC can be used in the following ways:

- MPIO should be used for host access. Multiple FC ports can be installed on the host SQL Server database to provide highly available connectivity to the LUNs accessible by the host.
- MPIO can be used on the host for providing resilient connections from host to storage. Similarly, MPIO can also be used in guest machines for resilient connections from the guest operating system to FC storage.

SQL Networking Considerations

This subsection discusses four key networking considerations for SQL Server: Dynamic Virtual Machine Queue, Single Root I/O Virtualization, IPSec, and QoS Bandwidth Management.

Dynamic Virtual Machine Queue

Virtual Machine Queue (VMQ) is a feature introduced in Windows Server 2008 R2 that has the Hyper-V server role installed and that uses VMQ-capable network hardware. VMQ uses hardware packet filtering to deliver packet data from an external virtual machine network directly to virtual machines, which reduces the overhead of routing packets and copying them from the management operating system to the virtual machine.

Windows Server 2012 introduces Dynamic Virtual Machine Queue (D-VMQ) for Hyper-V, which dynamically distributes the processing of incoming network traffic to host SQL Server processors, based on processor use and network load. In times of heavy network load, D-VMQ automatically uses more processors; when the network load decreases, D-VMQ relinquishes these processors.

Best Practices and Recommendations

Some Intel multicore processors may use Intel Hyper-Threading Technology. When Hyper-Threading Technology is enabled, the actual number of cores that are used by D-VMQ should be half the total number of logical processors that are available in the system. This is because D-VMQ spreads the processing across individual physical cores only, and it does not use hyper-threaded sibling cores.

As an example, if the machine has an Intel processor with four physical cores and Hyper-Threading Technology is enabled, it will show a total of eight logical processors. However, only four logical processors are available to VMQ. (VMQ will use cores 0, 2, 4, and 6.)

As shown in Figure 33, without the VMQ technology and RSS, the majority of network processing burdens CPU0 and ultimately limits the scale of the solution. With D-VMQs, processor cores are dynamically assigned to distribute the workload.

Figure 33: Dynamically distributed workload with D-VMQ for Hyper-V

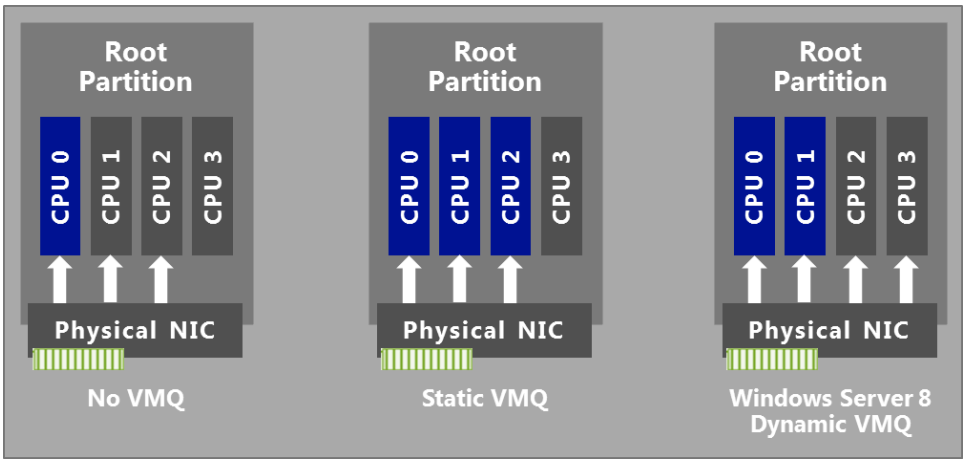

Single Root I/O Virtualization

The Single Root I/O Virtualization standard was introduced by the PCI-SIG, the special interest group that owns and manages PCI specifications as open industry standards. SR-IOV helps to virtualize demanding workloads like SQL Server that require high network and I/O performance. It does so by enabling virtual machines to perform I/O directly to the physical network adapter by bypassing the root partition. In Windows Server 2012, SR-IOV can be deployed in conjunction with key capabilities such as live migration to enable high network performance with availability.

SR-IOV provides extensions to PCI Express (PCIe) devices like network adapters to separate access to its resources among various PCIe hardware functions. Two of these functions are PCIe Physical Function (PF) and PCIe Virtual Functions (VFs):

- **PCIe Physical Function** is the primary function of the device and advertises its SR-IOV capabilities. The PF is associated with the Hyper-V parent partition in a virtualized environment.

- **PCIe Virtual Functions** are associated with the PF of the device. A VF shares one or more physical resources, such as memory and network ports, with the PF and other VFs on the device. Each VF is associated with a Hyper-V child partition in a virtualized environment.

Using Hyper-V Manager, you can enable SR-IOV in Windows Server 2012 when you create a virtual switch (Figure 34).[46]

Figure 34: Enabling SR-IOV in the Virtual Switch Properties window

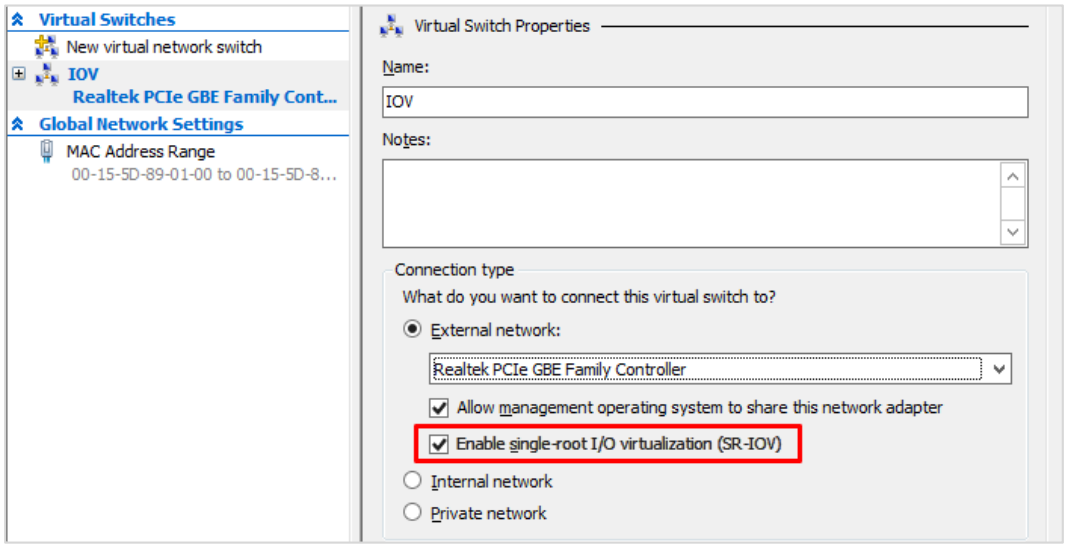

Once the virtual switch is created, SR-IOV should also be enabled while configuring a virtual machine in the Hardware Acceleration node (Figure 35).

Figure 35: Enabling the SR-IOV option from the Virtual Machine Properties window

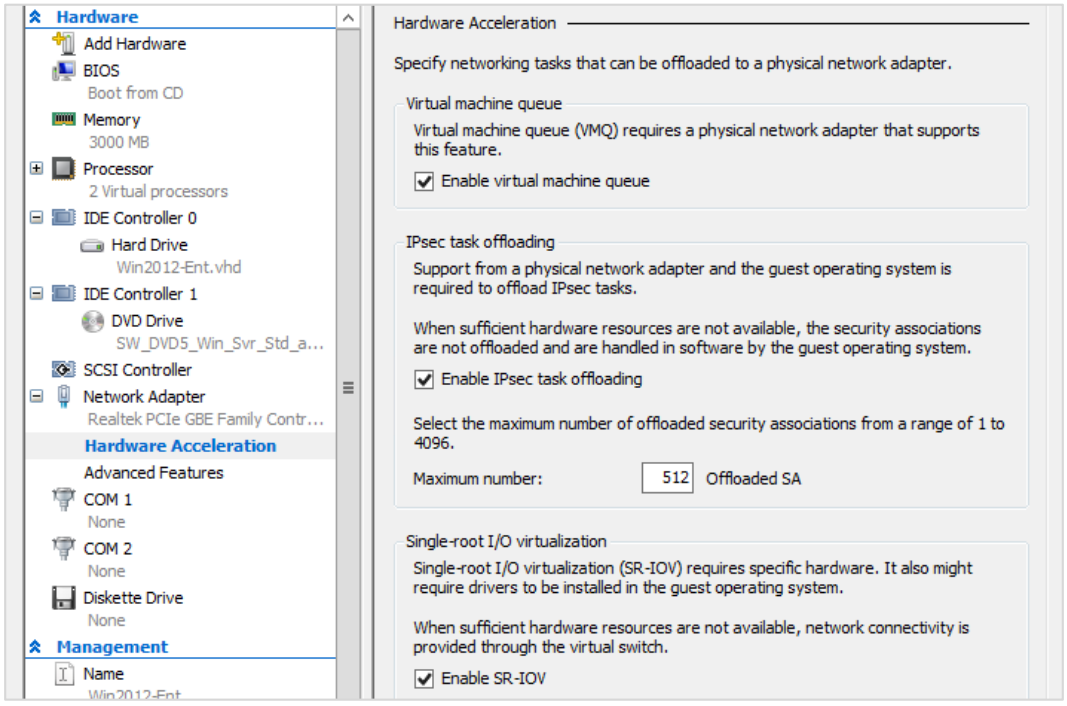

Figure 36 shows how SR-IOV attaches a physical NIC to a SQL Server virtual machine, which helps with the live migration of the virtual machine. In this process:

- SR-IOV is enabled for the SQL Server virtual machine, and its virtual function is assigned.
- A team is automatically created inside the SQL Server virtual machine, and the traffic flows down the virtual function (VF) path, not the software stack.
- When live migration is initialized, the team is broken and the VF is removed.
- At this point, live migration of the SQL Server virtual machine takes place from source to destination. (Traffic is now travelling via the synthetic software stack.)
- Upon arrival, the VF is reassigned and the team recreated.

Note that throughout this process, the SQL Server virtual machine always has connectivity.

Figure 36: SR-IOV support in Hyper-V

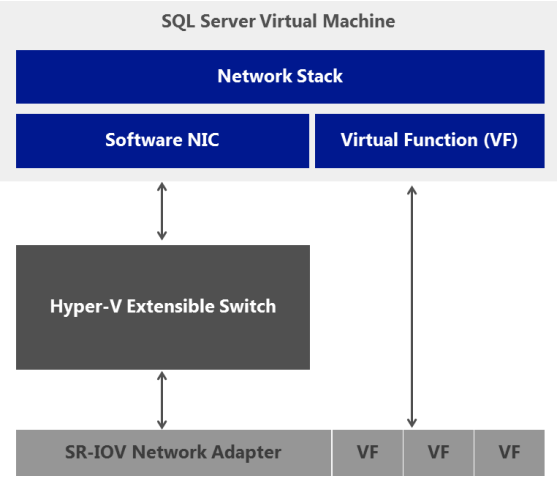

If the SQL Server virtual machine is configured to use SR-IOV, but the guest operating system does not support it, SR-IOV virtual functions are not allocated to the virtual machine. We recommend that you disable SR-IOV on all SQL Server virtual machines that run guest operating systems that do not support SR-IOV.[16]

Best Practices and Recommendations

SR-IOV provides the highest levels of networking performance for your virtualized SQL Server. Check with your hardware vendor for support because there may be a BIOS and firmware update required to enable SR-IOV.

IPsec

Internet Protocol Security (IPsec) Task Offload is a technology built into the Windows operating system that moves CPU workloads from the main computer's CPU to a dedicated processor on the network adapter. IPSec has been introduced to increase network performance and secure network traffic going to and from computers. It helps to preserve network traffic by using cryptographic security. As security requirements increase, IPSec allows a switch to higher security algorithms.

Cryptographic security services are implemented using IPsec, which protects communications over IP networks. IPsec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection. The Microsoft implementation of IPsec is based on Internet Engineering Task Force (IETF) standards.

IPsec is configured by using Windows Firewall with the Advanced Security dialog box in Windows 7, Windows Server 2008 R2, Windows Vista, and Windows Server 2008.

Best Practices and Recommendations

Use the cryptographic algorithms and methods that provide the minimum acceptable level of security while developing a security design. For example, if a 192-bit algorithm is capable of protecting network traffic, then do not use 256-bit or 384-bit algorithms.

QoS Bandwidth Management

Quality of Service is a prioritization technique that gives the ability to cost effectively manage network traffic and enhance user experiences in enterprise environments. QoS allows you to meet the service requirements of a workload or application in a SQL Server environment by measuring network bandwidth, detecting changing network conditions (such as congestion or availability of bandwidth), and prioritizing or throttling network traffic. QoS provides features like bandwidth management, classification and tagging, priority-based flow control, policy-based QoS, and Hyper-V QoS.[47]

QoS bandwidth management helps to set a throttling rate for a workload like SQL Server 2012. Both Minimum Bandwidth and Maximum Bandwidth enable organizations to enforce predictable network throughput for SQL Server 2012. Apart from bandwidth management, organizations can prioritize and tag traffic so that QoS is enforced from end-to-end across a data center.

Note that QoS Minimum Bandwidth and DCB are not designed to work together. Figure 37 shows that either Minimum Bandwidth or DCB on a computer running Hyper-V can be enabled, but not both technologies at the same time. However, you can enable both Minimum Bandwidth and DCB on one server running Windows Server 2012 for two separate networking stacks or NICs.[48]

Figure 37: Enabling Minimum Bandwidth and DCB on one server running Windows Server 2012 for two networking stacks

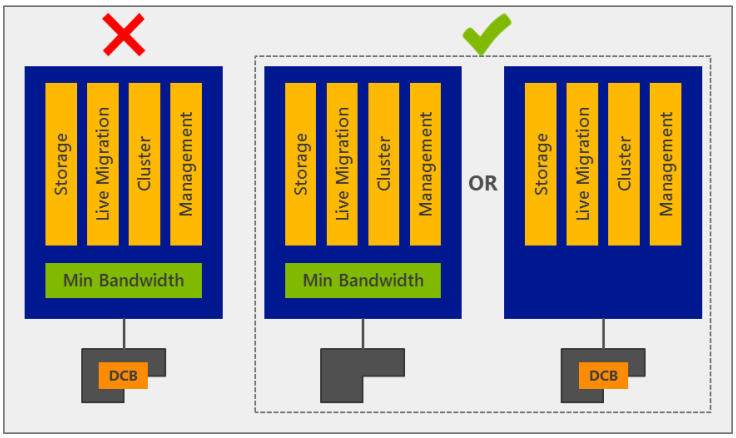

SQL Server Resiliency

Several features and technologies have been developed to help ensure the resiliency of SQL Server. This section discusses the following three: high availability with failover clustering, high availability with guest clustering, and AlwaysOn.

High Availability with Failover Clustering

The Hyper-V role in Windows Server 2012 helps to improve the efficiency of storage, compute, and network resources; it also helps to increase SQL Server availability in a virtualized environment. The Windows Server Failover Clustering (WSFC) feature in Windows Server 2012 enables creating and managing failover clusters. SQL Server AlwaysOn provides AlwaysOn failover cluster instances that use WSFC functionality to provide local high availability through redundancy at the server-instance level—a failover cluster instance (FCI). Working in conjunction with a Hyper-V failover cluster, FCI provides better availability data.

An FCI is a single instance of SQL Server that is installed across WSFC nodes and, possibly, multiple subnets. An FCI appears as an instance of SQL Server that provides failover from one WSFC node to another if the current node becomes unavailable. In Figure 38, the SQL Server instance is configured to be an FCI (instead of a standalone instance). The high availability of this instance is protected by the presence of redundant nodes in the FCI. When the SQL Server node fails, the client or application is transparently connected to the other SQL Server node in the cluster. This minimizes the downtime that the client or application experiences during a failure.[49]

Figure 38: Failover cluster instance with high availability protected

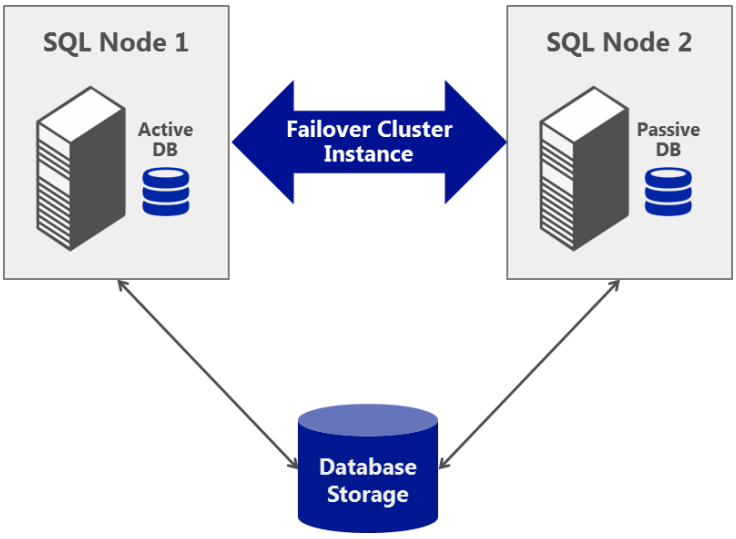

Best Practices and Recommendations

Create at least two virtual network switches in the Hyper-V Manager console. One is for the cluster heartbeat (private network switch), and the other is for cluster communication and iSCSI traffic (internal virtual switch).

Before the clusters are created, run a full validation test of cluster configuration using the Validate a Configuration Wizard in the Failover Cluster Manager or the Windows PowerShell Test-Cluster cmdlet.

In a production environment, use static IP addresses in conjunction with the virtual IP address of a failover cluster instance.

Do not couple standalone nodes together in an availability group if automatic failover is an important component of your high availability solution.

High Availability with Guest Clustering

Another high availability option for SQL Server involves guest clustering. Figure 39 illustrates guest clustering, which makes the provisioning of clusters possible in a virtual environment. Guest failover clusters are simply SQL Server failover clusters inside virtual machines, with nodes operating as virtual machines. The shared storage required to build a cluster is connected to each virtual machine by means of iSCSI. SQL Server guest cluster nodes can all run on the same physical host machine or on different physical host machines. Note that if SQL Server guest clustering is running on the same host, business operations and high availability will be compromised if the host becomes unavailable. Therefore, consider running guest cluster nodes of a SQL Server virtual machine on different physical Hyper-V host machines to maintain high availability and reduce interruptions in business operations.

Figure 39: Guest clustering for high availability in SQL Server

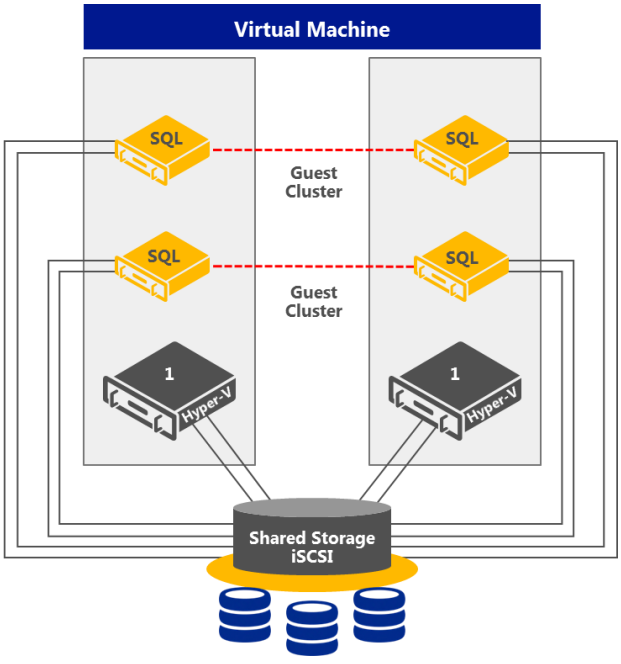

AlwaysOn

Increasingly, organizations require continuous uptime of critical applications, as well as protection of their critical data from planned and unplanned downtime. SQL Server 2012 hosted in the Hyper-V environment addresses availability and reliability of data at every level, while helping to keep costs and complexity low. For example, SQL Server 2012 AlwaysOn provides a high availability and disaster recovery solution that helps to protect against any data loss. It also helps to improve the availability of servers, applications, and databases so that the perceived downtime for users is minimized. AlwaysOn enables high availability using two key technologies: AlwaysOn Failover Cluster Instances and AlwaysOn Availability Groups.

AlwaysOn Failover Cluster Instances

AlwaysOn Failover Cluster Instances use Windows Server Failover Clustering to provide local high availability and redundancy at the FCI level. As discussed above, an FCI is a single instance of SQL Server that is installed across WSFC nodes and multiple subnets. FCI provides failover from one WSFC node to another when the current node fails. SQL Server 2012 takes advantage of WSFC functioning as a group of independent servers that work together to support AlwaysOn Availability Groups and SQL Server FCIs.

In WSFC, if a cluster node or service fails, the services that were hosted on that node can be automatically or manually transferred to another available node using the failover process. An FCI runs in a WSFC resource group containing one or more WSFC nodes. When the FCI is initiated, one of the nodes takes over the ownership of the resource group and brings its SQL Server instance online. In the event of a failover, the following steps are performed:

- All dirty pages in the buffer cache are written to disk, except in the case of a hardware or system failure.
- All respective SQL Server services in the resource group are stopped on the active node.
- The resource group ownership is transferred to another node in the FCI.
- The new resource group owner starts its SQL Server services.
- Client application connection requests are automatically directed to the new active node using the same virtual network name (VNN).

The Windows Server 2012 Failover Clustering Dynamic Quorum allows the AlwaysOn cluster to dynamically adjust the number of quorum votes required to keep running (Figure 40).[50, 51] Dynamic quorum management can simplify set up by as much as 80 percent. Further, it can help to increase the availability of a SQL Server cluster in failover scenarios in both virtualized and non-virtualized environments, with the ability to easily recalculate a quorum while maintaining a working cluster.

Figure 40: AlwaysOn FCI residing in a WSFS cluster

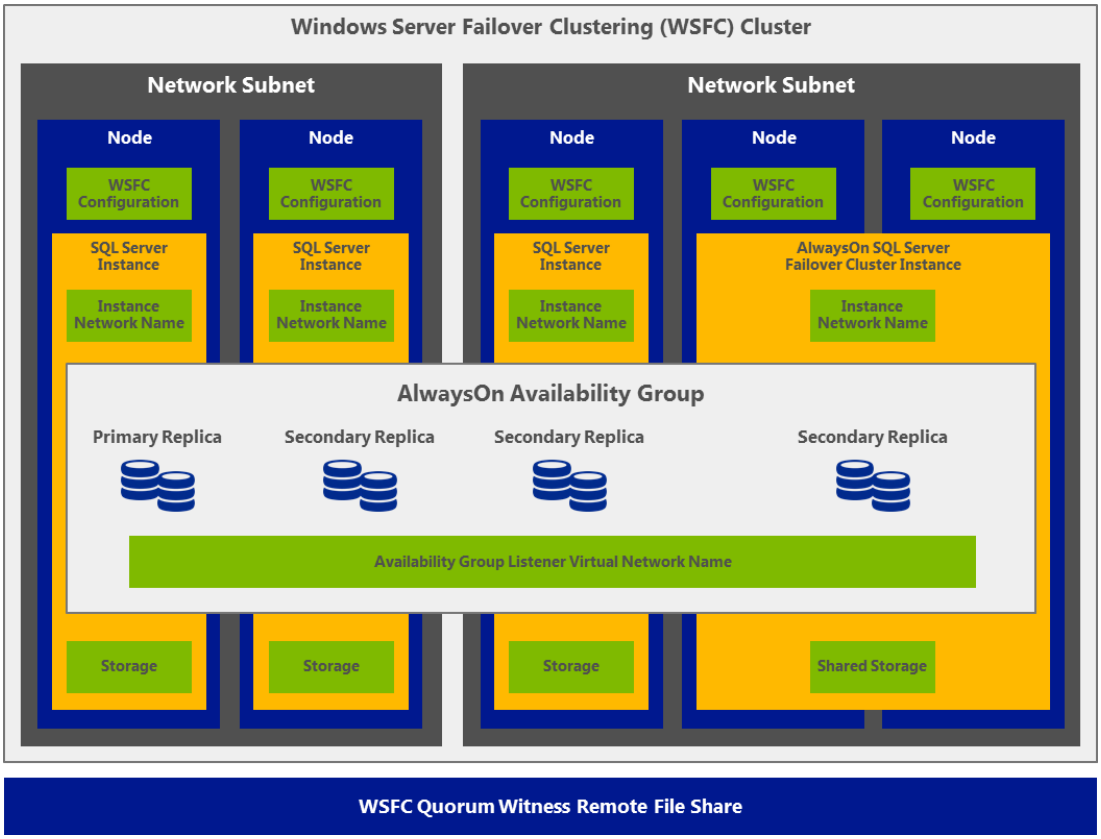

AlwaysOn Availability Groups

AlwaysOn Availability Groups are an enterprise-level high availability and disaster recovery solution introduced in SQL Server 2012. The solution is designed to maximize availability for one or more user databases. For AlwaysOn Availability Groups, the SQL Server instances need to reside on WSFC nodes. An Availability Group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. Secondary databases can be made available optionally for read-only access and some backup operations. Up to five availability replicas (one primary and four secondary) can be supported. An availability replica is an instance of an Availability Group that is hosted by a specific instance of SQL Server and maintains a local copy of each availability database. Availability Groups support several forms of failover: automatic, planned manual, and forced manual.

Availability Modes

With AlwaysOn Availability Groups, alternate availability modes can be used depending on the placement of servers—remote or local:

- **Asynchronous-commit mode**: This availability mode is a disaster recovery solution that works well when the availability replicas are distributed over considerable distances.
- **Synchronous-commit mode**: This availability mode emphasizes high availability and data protection over performance, at the cost of increased transaction latency. A given Availability Group can support up to three synchronous-commit availability replicas, including the current primary replica.

Availability Replicas

An Availability Group works on a distinctly separate set of user databases, known as *availability databases*, that fail over together. Each of these availability databases is hosted by an availability replica. As noted above, there are two types of availability replicas: one primary replica that hosts the primary databases and one to four secondary replicas that host a set of secondary databases. The availability replicas serve as potential failover targets for the Availability Group. An availability replica provides redundancy only at the database level—for the set of databases in one Availability Group. Failovers are not caused by database issues (for example, a database becoming suspect due to loss of a data file or corruption of a transaction log).

The primary replica makes the primary databases available for read-write connections from clients using data synchronization that occurs at the database level. The primary replica sends transaction log records of each primary database to every secondary database. Every secondary replica caches the transaction log records ("hardens the log") and then applies them to its corresponding secondary database. Data synchronization occurs between the primary database and each connected secondary database, independently of the other databases. Therefore, a primary or secondary database failure does not affect its corresponding primary or secondary databases.

Deploying AlwaysOn Availability Groups requires a 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.

A WSFC resource group is assigned to each Availability Group that is created (Figure 41).[52] The WSFC cluster monitors this resource group to evaluate the health of the primary replica. The quorum for AlwaysOn Availability Groups is based on all nodes in the WSFC cluster, regardless of whether a given cluster node hosts any availability replicas.

Figure 41: Availability Group residing in WSGC resource group

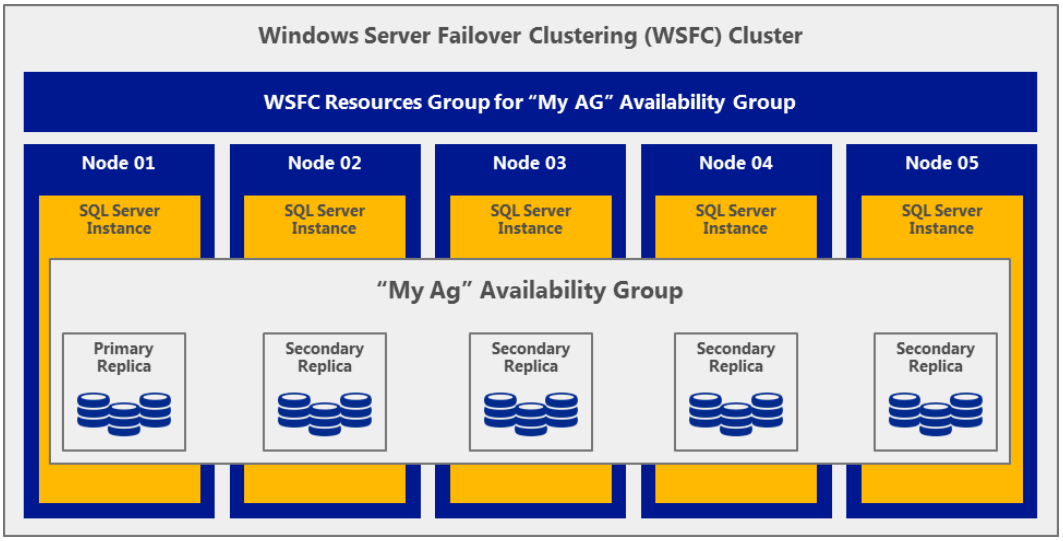

System Center 2012 SP1

System Center 2012 SP1 provides several components that give IT the ability to streamline infrastructure management and—as discussed in this guide specifically—to better deploy, manage, maintain, and protect SQL Server in a virtualized environment. This section, organized as follows, explores these components in greater detail:

- Comprehensive SQL Server Management
- Virtual Machine Manager
- App Controller
- Service Delivery and Automation (Service Manager, Orchestrator, and Cloud Services Process Pack)
- Operations Manager
- Data Protection Manager

Comprehensive SQL Server Management

Cloud computing is transforming the way organizations provide and consume IT services with the promise of more productive infrastructure and more predictable applications. System Center 2012 SP1 delivers on this promise by enabling your enterprise to benefit from private, hosted, and public cloud computing while still supporting your unique business needs. It helps to organize your IT assets—network, storage, and compute—into a hybrid cloud model spanning private cloud and public cloud services from a single console view.

Infrastructure Management: System Center 2012 SP1 provides a common management toolset to help you configure, provision, monitor, and operate your IT infrastructure. If your infrastructure is like that of most organizations, you have physical and virtual resources running heterogeneous operating systems. The integrated physical, virtual, private, and public cloud management capabilities in System Center 2012 SP1 can help you ensure efficient IT management and optimized ROI of those resources.

Service Delivery and Automation: System Center 2012 SP1 helps you simplify and standardize your data center with flexible service delivery and automation. Using the Service Manager and Orchestrator components of System Center 2012, you can automate core organizational process workflows like incident management, problem management, change management, and release management. You can also integrate and extend your existing toolsets and build flexible workflows (or runbooks) to automate processes across your IT assets and organizations.

Application Management: System Center 2012 SP1 offers unique application management capabilities that can help you deliver agile, predictable application services. Using the App Controller, Operations Manager, and Virtual Machine Manager components of System Center 2012 SP1, you can provide *Applications as a Service*—where a "service" is a deployed instance of a cloud-style application along with its associated configuration and virtual infrastructure.

Virtual Machine Manager

System Center 2012 SP1 Virtual Machine Manager (VMM) is the nerve center of SQL Server deployment. This subsection takes a deeper dive into VMM by discussing the following topics: centralized fabric configuration, virtual machine creation, virtual machine deployment, Dynamic Optimization, virtual machine priority and affinity, availability sets, and the private cloud/SQL cloud.

Centralized Fabric Configuration

As shown in Figure 42, VMM enables IT administrators to easily configure and manage virtualization hosts, networking, and storage resources in order to rapidly create and deploy SQL Server virtual machines and services.

Figure 42: Overview of System Center 2012 SP1 Virtual Machine Manager

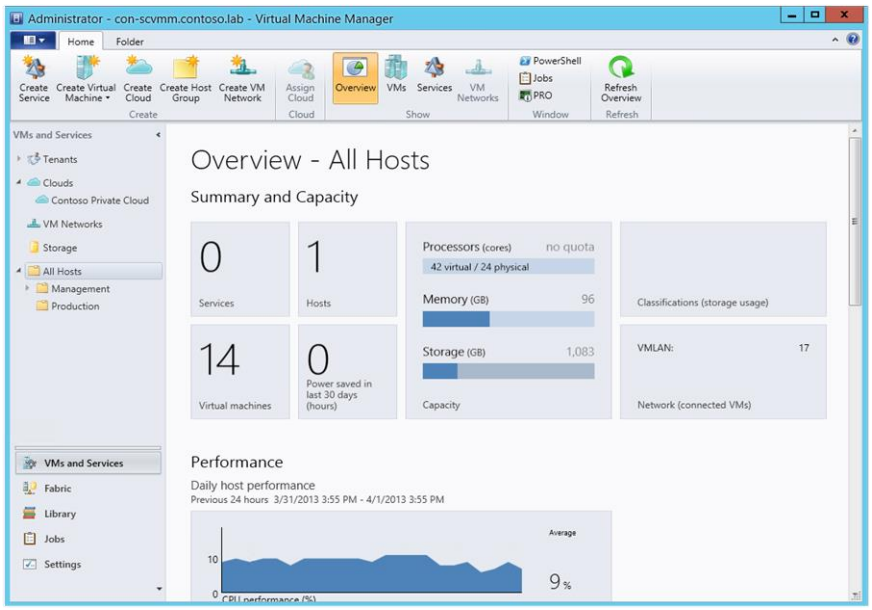

In addition to managing Hyper-V, VMM provides management for Citrix XenServer and VMware ESX/i hosts and host clusters. To help organize hosts and the deployment of virtual machines, IT administrators can create host groups based on considerations such as physical site location or resource allocation.

For networking, VMM manages resources such as logical networks, IP address pools, and load balancers that are used to deploy virtual machines and services. VMM also manages storage resources (such as storage classifications, LUNs, and storage pools) that are made available to Hyper-V hosts and host clusters.

Virtual Machine Creation

The VMM Management console provides several capabilities and features that can be used to accelerate and optimize deployment of SQL Server in a virtualized environment.

Physical-to-Virtual Conversions

First, VMM offers an inbox Physical-to-Virtual (P2V) capability to quickly and efficiently convert physical SQL Servers into virtual SQL Servers to run on Hyper-V. VMM offers two methods for conversion of physical machines: online and offline:

- **Online Conversion**: With an online conversion, the source computer continues to perform normal operations and is available throughout the process. VMM creates a copy of local NTFS volumes and data for VSS-aware applications. VMM uses the Volume Shadow Copy Service (VSS) to ensure that data is backed up consistently while the server continues to service user requests. VMM then uses this read-only snapshot to create a VHD.
- **Offline Conversion**: For a busy SQL Server, the point-in-time capture of the local copy for an online P2V would be out-of-date very quickly, so an automated offline conversion may be more appropriate. Here, the source computer restarts in the Windows Preinstallation Environment (Windows PE), and then VMM clones the volume to a VHD. Offline P2V conversion is the only method to reliably migrate FAT volumes, and it is the recommended method for converting domain controllers. Offline P2V conversion often is the most reliable way to ensure data consistency, especially in mission-critical scenarios like converting demanding SQL Server environments.

Virtual Machine Profiles and Templates

In VMM, a *profile* is a library resource containing specifications that can be applied to a new virtual machine or a virtual machine template. *Templates* encapsulate a standard set of configuration settings that can be used when creating a virtual machine. Templates can help you to quickly create virtual machines with consistent hardware and operating system settings. This can be extremely useful for the rapid deployment of SQL Server virtual machines into an infrastructure. Templates can also be used to restrict the virtual machine settings available to self-service users creating new virtual machines.

Core Profiles

Profiles are used when creating templates. A template typically consists of a hardware profile, an operating system profile, and a VHD that will be used by the virtual machine. The VHD might be stored in the VMM library, or it might be a disk from an existing virtual machine.

- **Hardware Profile**: A hardware profile defines hardware configuration settings such as CPU, memory, network adapters, video adapter, DVD drive, floppy drive, COM ports, and the priority given the virtual machine when allocating resources on a virtual machine host (Figure 43).

Figure 43: Hardware profile in Virtual Machine Manager

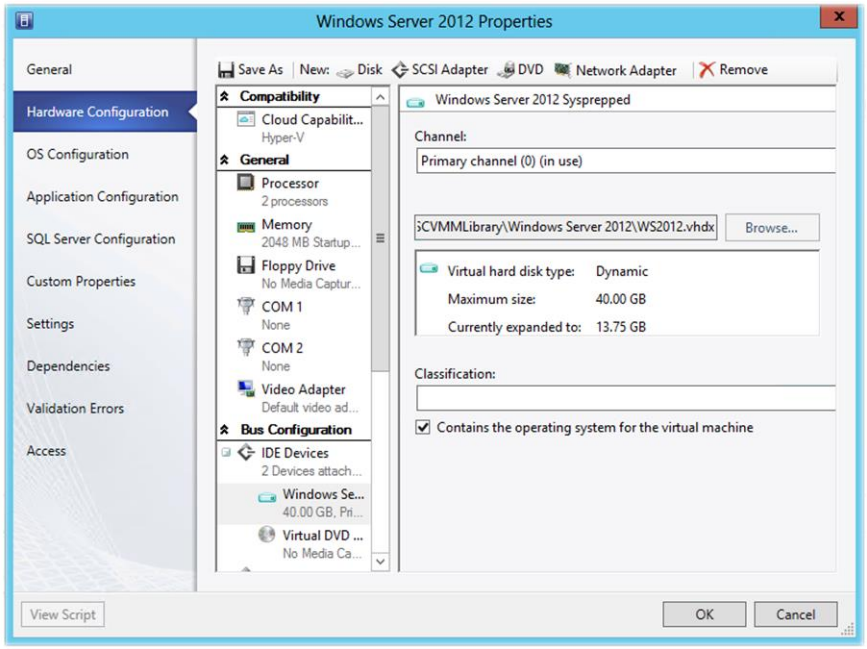

- **Guest Operating System Profile**: A guest operating system profile defines operating system-configured settings, which are applied to a virtual machine created from a template (Figure 44). This profile defines common operating system settings, such as type of operating system, roles and features to be enabled inside the operating system, computer name, administrator password, domain name, product key, time zone, answer file, and run-once file.

Figure 44: Guest operating system profile with roles and features displayed

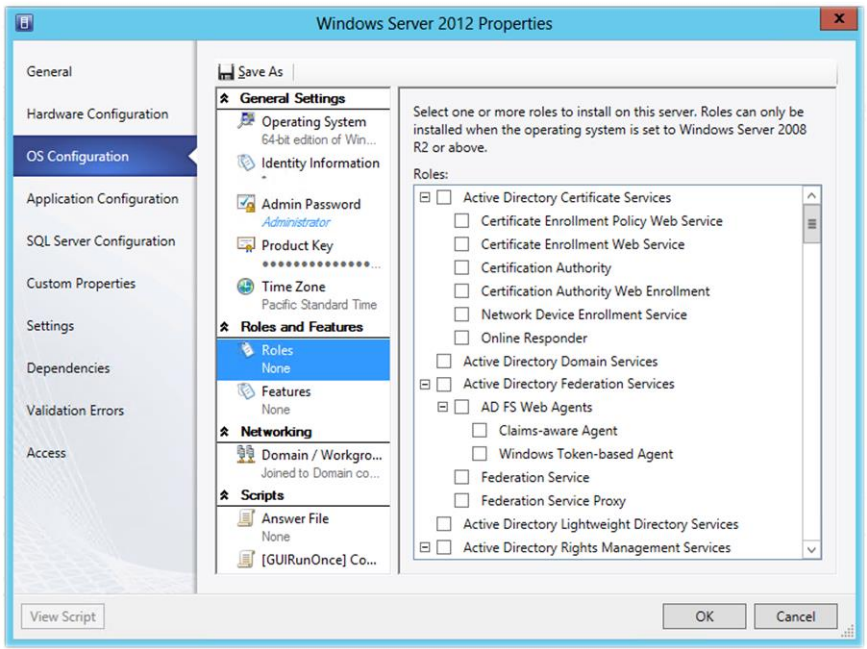

These profiles on their own, however, are not enough to be classified as a template. A template contains several other key elements that help to accelerate the deployment of new virtual machines into the

infrastructure. Templates are database objects stored in the library catalog of the VMM database. They are not represented by physical configuration files. Templates can be created in two ways:

- From an existing virtual hard disk or template stored in the library.
- From an existing virtual machine deployed on a host.

Best Practices and Recommendations

One of the easiest ways to create a template is to generate a new, blank virtual machine with your desired hardware settings; install the Windows operating system of your choice; install relevant updates or patches; and, once complete, shut the virtual machine down. You can then use the VMM Template Creation Wizard to transform this "gold virtual machine" into a new template. As part of the process, VMM automatically syspreps the guest operating system to ensure its uniqueness in future deployments. If you want to keep the original, take a clone of the gold virtual machine before the template creation process. Once VMM has finished creating the template, it stores the relevant files in the library. You can then begin deploying new virtual machines from this template.

Other Key Profiles

After generating a sysprepped template, you can start to use some other profiles to enhance the template and accelerate deployment of virtual machines that have specific applications within them. These profiles include the application profile and SQL Server profile:

- **Application Profile**: An application profile provides instructions for installing Microsoft Web Deploy applications, Microsoft Server Application Virtualization applications, and Microsoft SQL Server data-tier applications. This profile also provides instructions for running scripts when deploying a virtual machine as part of a service (Figure 45).

Figure 45: Application profile for SQL Server

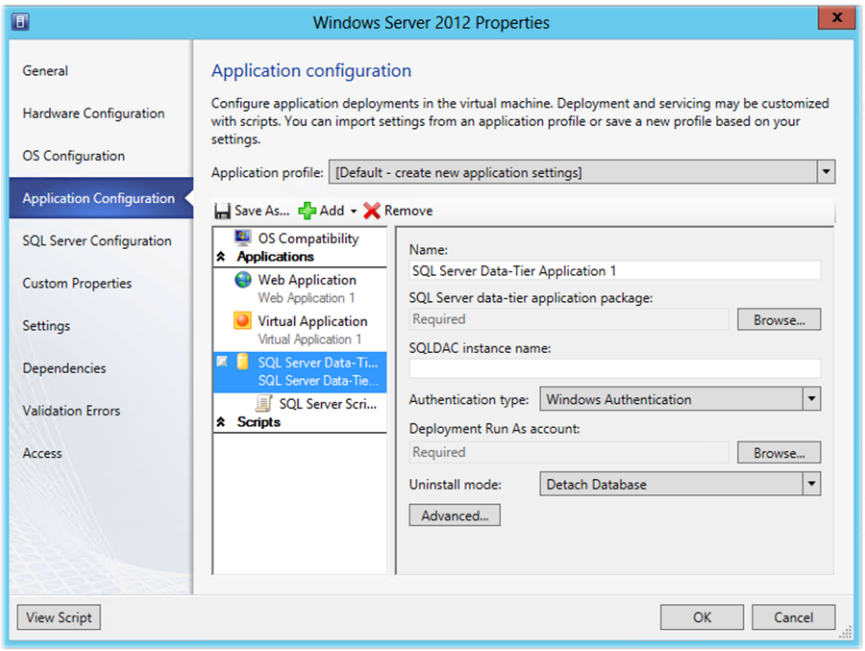

- **SQL Server Configuration Profile**: A SQL Server configuration profile provides instructions for customizing an instance of SQL Server for a SQL Server DAC when deploying a virtual machine as part of a service (Figure 46).

Figure 46: SQL Server configuration profile

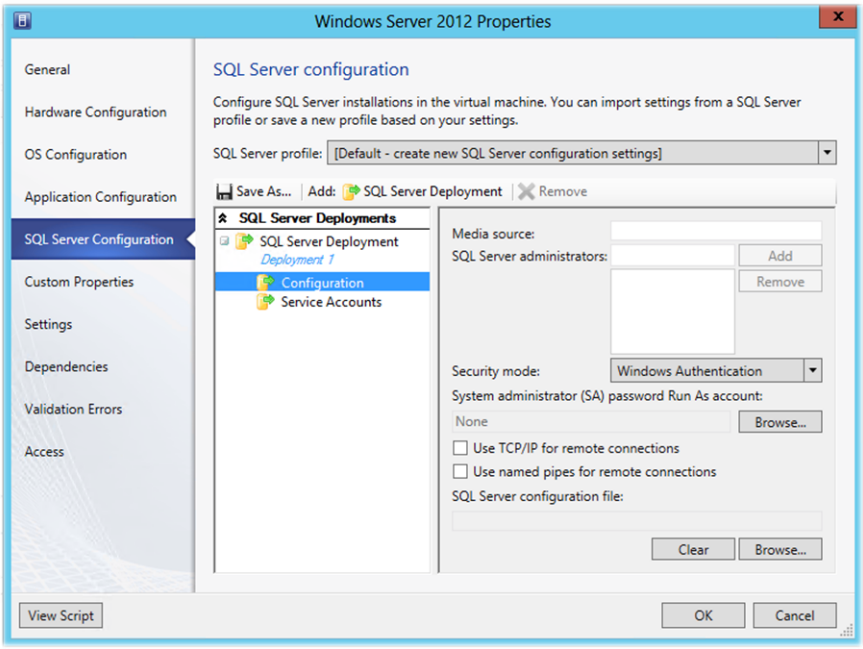

As shown in Figure 45, you can configure an application profile that can be used to accelerate deployment of virtual machines that contain certain key characteristics:

- **Web Deploy**: The first characteristic, websites, makes use of Microsoft Web Deploy capabilities to accelerate deployment of virtual machines that are to host websites and provide other specific web functionality.
- **Server App-V**: The second characteristic is application virtualization. You can use Microsoft Server Application Virtualization (Server App-V) to create virtual application packages. Virtual application packages are images of applications that can be copied to a computer running the Server App-V Agent and started without requiring a local installation. The application then runs as if it is a locally installed application.

  Server App-V builds on the technology used with Microsoft Application Virtualization (App-V) by separating the application configuration and state from the underlying operating system running on computers in a data center environment. Server App-V allows for dynamic composition of application and hardware images, which can help to significantly reduce the number of images that need to be managed. Server App-V also enables automation of deployment and management scenarios, which can improve the reliability, availability, and serviceability of data center applications.

  Note that not all applications are supported for use with Server App-V. For example, applications such as antivirus software that require device or kernel driver support are not supported. Server App-V is primarily designed for use with business applications or the business tiers of multi-tiered applications. Consequently, some large server applications such as Microsoft Exchange Server, Microsoft SQL Server, and Microsoft SharePoint are not supported.

- **SQL Server Data-Tier Applications**: The third characteristic is integration with SQL Server Data-Tier Application packages. A data-tier application (DAC) is a logical database management entity that defines all SQL Server objects—like tables, views, and instance objects (including logins)—associated with a user's database. A DAC is a self-contained unit of SQL Server database deployment that enables data-tier developers and database administrators to package SQL Server objects into a portable artifact called a *DAC package* (DACPAC). Ultimately, the DACPAC helps to accelerate the deployment of SQL Server.

  To take advantage of DAC functionality, you need an installation of SQL Server. As discussed earlier, you can use VMM to sysprep a gold virtual machine, which contains an updated version of a particular Windows Server operating system. Thanks to innovation in SQL Server deployment capabilities, you can take this approach one step further and include a sysprepped image of SQL Server as part of the gold virtual machine. This helps to ensure that when you deploy new virtual machines with the intention of using SQL Server within them, there is little or no extra configuration required to use the SQL services.

Gold SQL Virtual Machine

To create a gold SQL virtual machine, start with an initial gold virtual machine. You can create a new gold virtual machine and apply updates to it, or you can reuse a template created earlier and deploy a new virtual machine from it. Once the virtual machine is deployed, create a directory on the local drive within the guest operating system, which will hold the SQL installation files. Copy the SQL installation files to this directory, but note that you only need to copy the files of the architecture you are going to use (x86 or x64). Once the files are copied, run the set-up process, selecting advanced options from the Installation Center. For example, you can walk through the steps needed to install SQL Server in a sysprepped manner, or you can export the SQL configuration file for use by VMM when deploying SQL Server at a later date.

Once the process is complete, VMM can be used to sysprep the entire operating system image, creating a new template that VMM stores in its library. This means that there are now two gold virtual machines:

- Gold virtual machine for deployment of a standard Windows Server.
- Gold SQL virtual machine for deployment of virtual machines that contain SQL Server.

To complete the template, you can edit the SQL Server profile to include relevant information about the SQL Server installation now encased within the VHD associated with the template. This includes details such as specifying the location of the installation files inside the guest operating system and identifying SQL administrators or service accounts.

Service Templates

Regular templates, like the standard gold virtual machine, can be deployed as they are, without further modification. However, a SQL Server-configured template, like the gold SQL virtual machine, must be deployed as part of a service. To deploy services, *service templates* are required.

In VMM, a *service* is a set of virtual machines configured and deployed together and managed as a single entity (for example, a deployment of a multi-tier line-of-business application). In the VMM Management console, you can use the Service Template Designer to create a service template that defines the configuration of the service (Figure 47). A service template includes the following information:

- Virtual machines that are deployed as part of the service.
- Applications to be installed on the virtual machines.
- Networking configuration needed for the service (including the use of a load balancer, if required).

Figure 47: Service Template Designer with SQL tier on the right

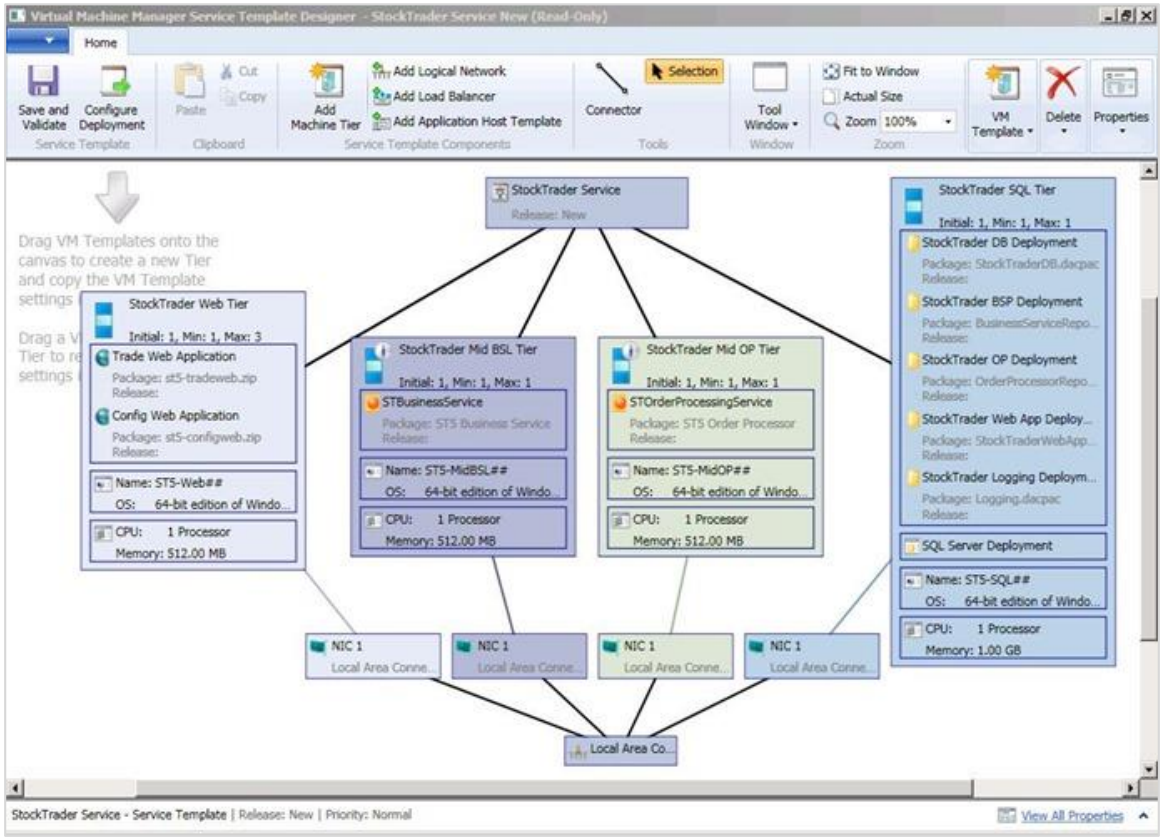

After the service template is created, you can deploy the service to a private cloud or to virtual machine hosts. Once the service is deployed, you can update the service template and apply those changes to the existing service. Or, you can deploy more virtual machines to the existing service in order to provide additional resources for it.

Best Practices and Recommendations

If you want to deploy single virtual machines containing SQL Server rather than a multi-tiered service, VMM provides a wizard that can help. A single-tier service template is easy to configure and requires little or no modification from the gold SQL virtual machine. Once the template is saved and validated, administrators can quickly and efficiently deploy new SQL Server virtual machines.

Virtual Machine Deployment

A virtual machine template can accelerate deployment of SQL Server virtual machines, mainly due to the ability of VMM to encase all necessary files and options within the template. When deploying the template itself, VMM also provides guidance on placement and moves the template around the infrastructure if it detects that better performance for that workload can be delivered from a different position.

Intelligent Placement

VMM can help to identify the most appropriate physical host servers for virtualized SQL Server workloads (Figure 48). Called *Intelligent Placement*, this technology not only can make administrative tasks easier, but also can help to ensure that data center resources are deployed properly and align with business goals.

Figure 48: Intelligent Placement in VMM

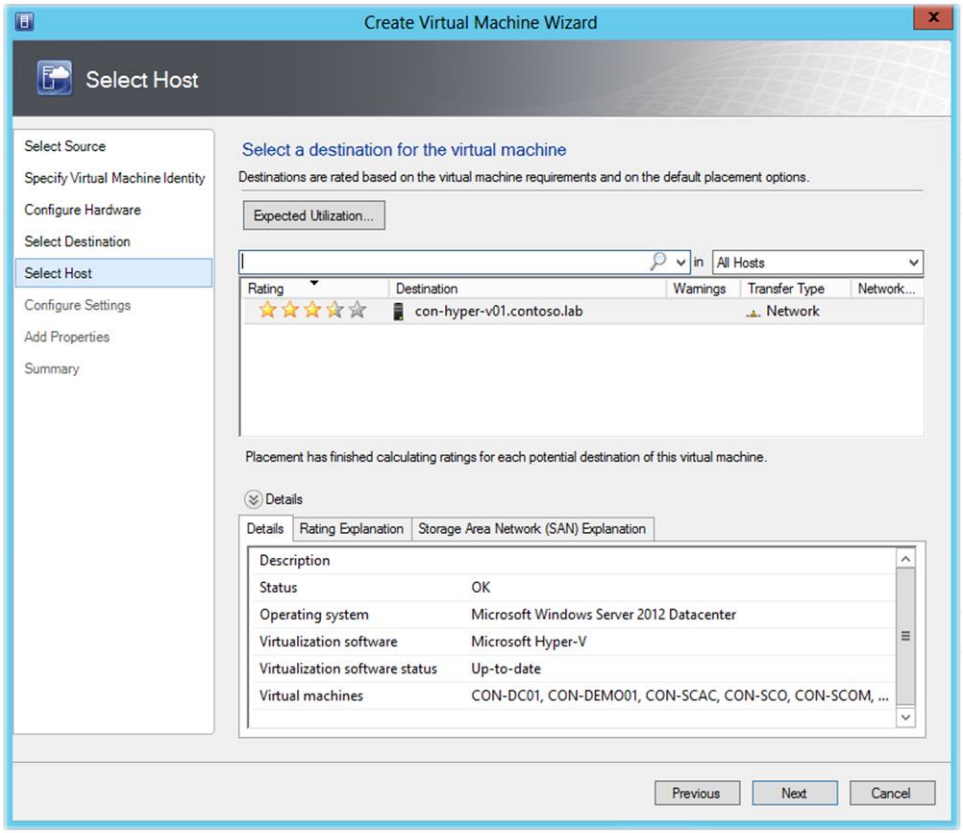

Intelligent Placement in VMM inputs host system data, workload performance history, and administrator-defined business requirements into sophisticated algorithms. This provides easy-to-understand, ranked results that can take the guesswork out of the placement task and help to ensure that workloads are spread across physical resources for optimum performance.

This placement also takes into account situations where a virtual machine requires specific hardware offload capabilities, such as SR-IOV, as defined as part of the template. If these are not available on a particular host, that host will not receive a star ranking as part of the Intelligent Placement destinations.

Storage Classification

VMM also provides the ability for an administrator to apply simple classifications to storage, which can be used for storing and running SQL Server virtual machines. Storage can be classified in any way that the administrator wants, but common examples include terms such as *Bronze*, *Silver*, and *Gold*, which may represent I/O characteristics, performance, and redundancy of the underlying storage array. For example, Bronze could be slower SATA drives in an older SAN, Silver could be SAS drives in a newer array, and Gold could be solid-state drive storage. These storage classifications can be used in the SQL virtual machine templates so that VMM automatically ensures that the chosen type of storage will be used for a particular deployment.

Dynamic Optimization

Once SQL Server virtual machines have been deployed onto the Hyper-V cluster, VMM actively monitors key cluster and host metrics, such as CPU, Memory, Disk, and Network, to see if it can better balance the virtual machine workloads across different hosts (Figure 49). For example, you may have several hosts in a cluster, but one of the hosts has some SQL virtual machines that are exhibiting higher levels of demand than others on other hosts. Through Dynamic Optimization, VMM can recognize this and automatically live migrate, with no downtime, some of the other virtual machines on the busy host to less-busy hosts, freeing up valuable resources to be used by the demanding SQL virtual machines. This helps to ensure that workloads such as SQL Server inside the virtual machines always receive the resources they need to meet demand, without impacting other workloads running on the cluster.

Figure 49: Dynamic Optimization in VMM

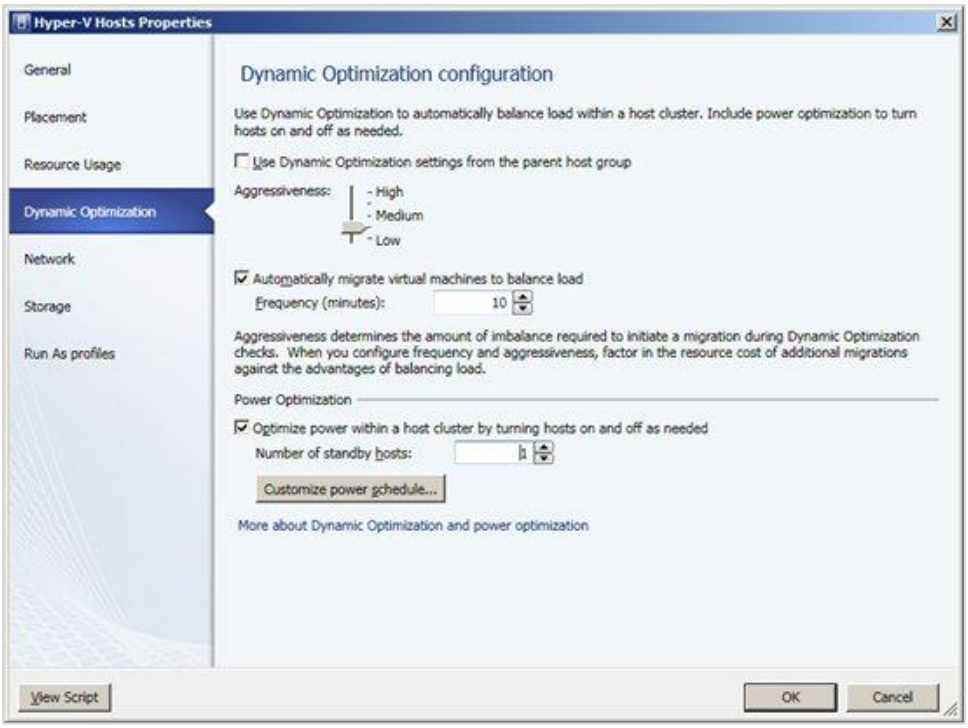

Best Practices and Recommendations

Dynamic Optimization can be configured on a host group to migrate virtual machines within host clusters with a specified frequency and aggressiveness. Aggressiveness determines the amount of load imbalance that is required to initiate a migration during Dynamic Optimization. By default, virtual machines are migrated every 10 minutes with medium aggressiveness. When configuring frequency and aggressiveness for Dynamic Optimization, an administrator should factor the resource cost of additional migrations against the advantages of balancing load among hosts in a host cluster. By default, a host group inherits Dynamic Optimization settings from its parent host group.

Virtual Machine Priority and Affinity

If you deploy virtual machines on a host cluster, you can use VMM to configure priority settings for them. With these settings, the cluster starts high-priority virtual machines before medium-priority or low-priority virtual machines. This ensures that high-priority virtual machines, like those running SQL Server, are allocated memory and other resources first for better performance. Also, after a node failure, if the high-priority virtual machines do not have the necessary memory and other resources to start, the lower priority virtual machines will be taken offline to free up the necessary resources. Virtual machines that are preempted are later restarted in priority order.

You can also configure the virtual machine priority setting in a virtual machine template, so that any virtual machines created with that template will have the specified virtual machine priority.

VMM also provides the ability for the administrator to influence the placement of virtual machines on the nodes of the host cluster by defining *preferred owners* and *possible owners* for the virtual machines (Figure 50). This helps to ensure that certain virtual machines can only run on certain hosts or that certain virtual machines will never run on a particular host.

Figure 50: Defining preferred and possible owners in VMM

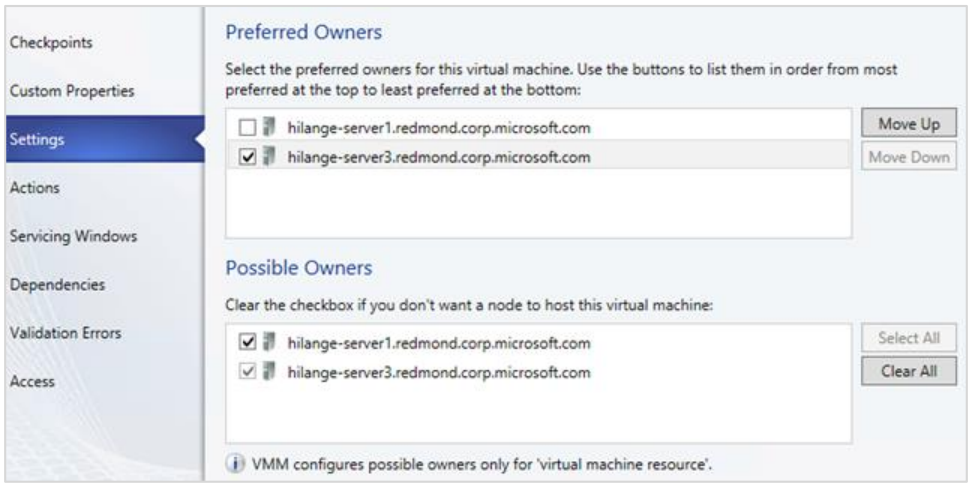

Availability Sets

When you place multiple virtual machines in an availability set, VMM attempts to keep those virtual machines on separate hosts whenever possible (Figure 51). Using the Availability Sets setting helps to improve continuity of service. Another way to configure this setting is to use Windows PowerShell commands for failover clustering. In this context, the setting appears in the *Get-ClusterGroup* listing and is called *AntiAffinityClassNames*. Note that you can also configure availability sets in a service template to specify how virtual machines created with that template should be placed on hosts.

Figure 51: Availability set for SQL virtual machines

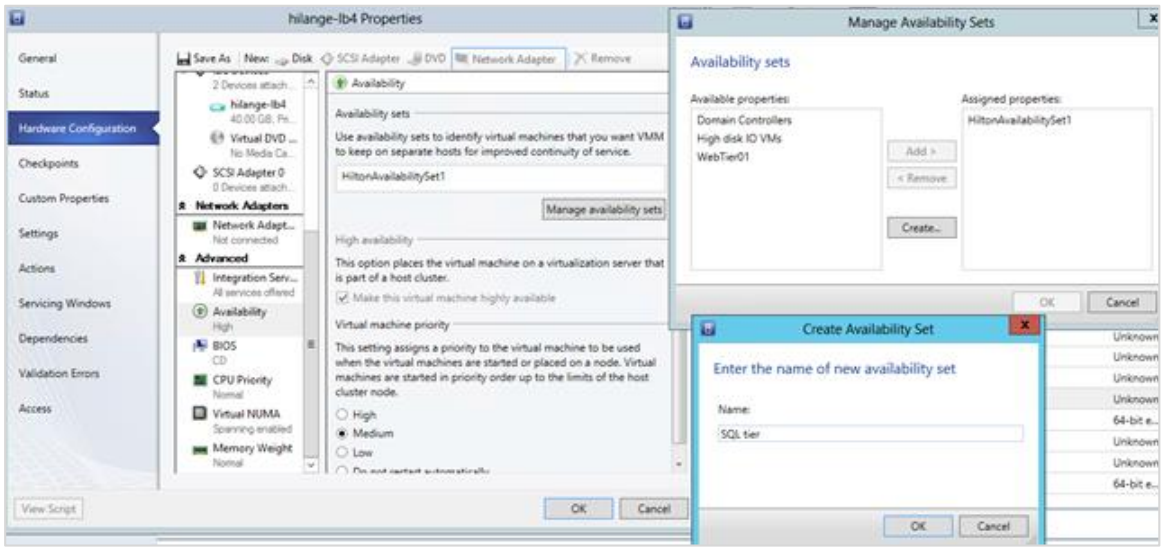

Best Practices and Recommendations

When creating a guest cluster (such as an AlwaysOn Availability Group) virtualized on top of a Hyper-V host cluster, it can be beneficial to keep the individual SQL Server nodes on separate hosts. If one physical host is lost, it takes down only a single node of the of AlwaysOn Availability Group because the availability set within VMM ensures that the AlwaysOn Availability Group nodes are running on separate hosts in the Hyper-V cluster.

Private Cloud/SQL Cloud

A private cloud is provisioned and managed on-premises by an organization. This cloud is deployed using the organization's own hardware to capitalize on the advantages of the private cloud model. Through VMM, an organization can quickly and easily manage the private cloud definition, access to the private cloud, and the underlying physical resources (Figure 52). VMM also provides granular, role-based access to end users, application owners, or database administrators.

Figure 52: Create Cloud Wizard in VMM

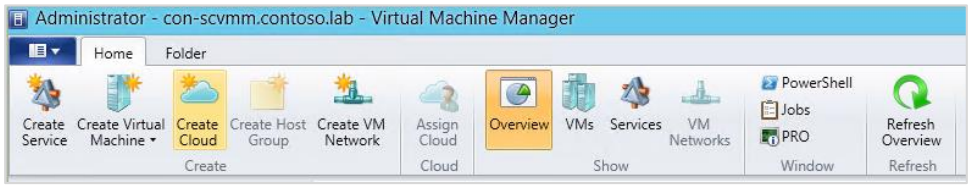

Through VMM, a private cloud provides the following benefits:

- **Resource Pooling**: Through the private cloud, administrators can collect and present an aggregate set of resources, such as storage and networking resources. Resource usage is limited by the capacity of the private cloud and by user role quotas.
- **Opacity**: Self-service users have no knowledge of the underlying physical resources.
- **Self-Service**: Administrators can delegate management and use of the private cloud while retaining the opaque usage model. Self-service users do not need to ask the private cloud provider for administrative changes beyond increasing capacity and quotas as their needs change.
- **Elasticity**: Administrators can add resources to a private cloud to increase capacity.
- **Optimization**: Usage of underlying resources is continually optimized without affecting the overall private cloud user experience.

When creating a private cloud, you select the underlying fabric resources that will be available, configure library paths for users, and set the capacity. Therefore, before you create a private cloud, you should configure the fabric resources, such as storage, networking, library servers and shares, host groups, and hosts.

Best Practices and Recommendations

From a SQL Server perspective, an IT administrator can define a cloud that can be used exclusively with SQL Server virtual machines. The administrator defines the capacity of the cloud, and the cloud uses elements such as the storage classifications discussed earlier to ensure that all virtual machines placed in it use a certain tier of storage. In addition, certain virtual machine templates and service templates can be assigned to the SQL cloud. This ensures that the only virtual machines that can be deployed into this cloud are those that already have SQL Server within the templates, thereby optimizing deployment.

Once the SQL cloud is created, you can assign access to certain users and groups within the IT infrastructure, such as database administrators. Through rich, granular role-based controls, you can delegate who can see what inside the cloud, as well as who can perform which tasks associated with it (Figure 53). Users who are part of the newly created group can access the cloud and associated virtual machines, templates, and service templates through the VMM Management console or, for a true self-service experience, through System Center 2012 SP1 App Controller.

Figure 53: Configuring user role actions in VMM

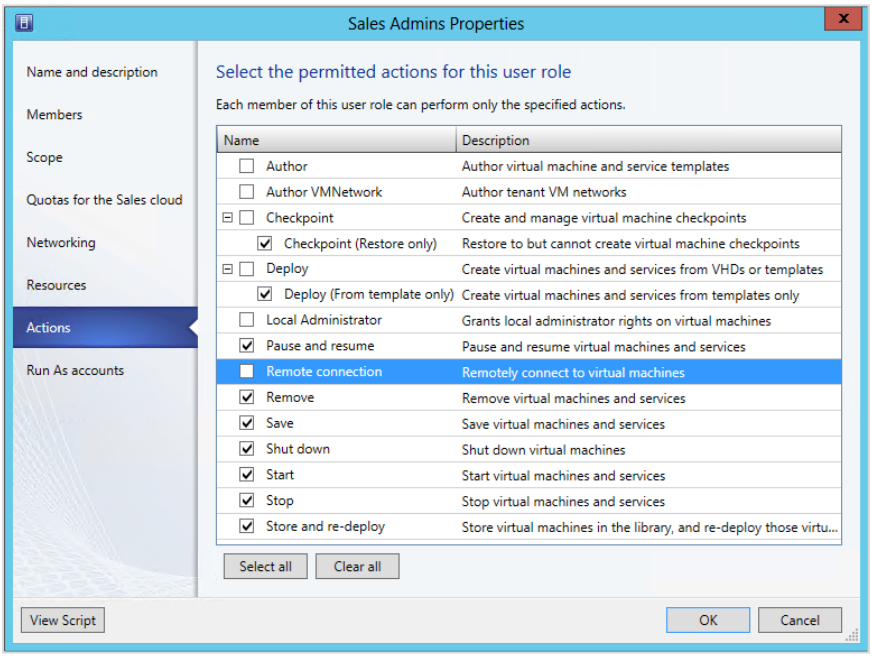

App Controller

Among the advantages of a private cloud is the ability to quickly provision and deprovision compute, networking, and storage resources through virtual machines. With System Center 2012 SP1 App Controller, IT administrators in your organization can give certain users (such as database administrators) the ability to access and consume private and public cloud infrastructure by self-provisioning standardized virtual machines in a controlled environment. This helps to reduce administrative overhead and improve time to market.

The following example shows the steps a database administrator (DBA) might take to self-provision a SQL Server virtual machine using App Controller. When the DBA first logs on to the App Controller interface, an overview screen is presented (Figure 54). This screen is dynamically generated based on identity; therefore, it automatically displays what the DBA can access. By selecting the *Clouds* option on the left side of the screen, the DBA has a visible representation of accessible clouds. In this example, note that the DBA has access not only to the Contoso SQL Cloud, but also to the Contoso Public Cloud.

Figure 54: Using App Controller to access a SQL private cloud

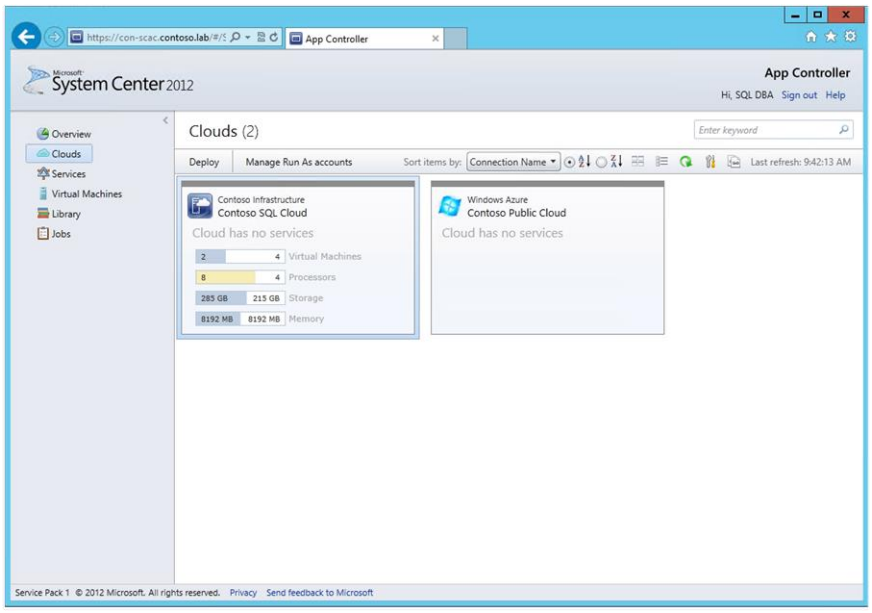

In Figure 55, the database administrator selects *Virtual Machines* on the left side of the screen. This brings up the list of current virtual machines visible to this DBA. What is important to note here is that the DBA only sees virtual machines that IT administrators have specifically provided and enabled for consumption. The rest of the virtual machines on a particular host or cluster are not visible, even though they may be running there. In addition, the DBA is only able to perform certain tasks on the virtual machines. In this example, the DBA can start, stop, shut down, and deploy new virtual machines, but cannot pause or save them.

Figure 55: Virtual Machine view in App Controller

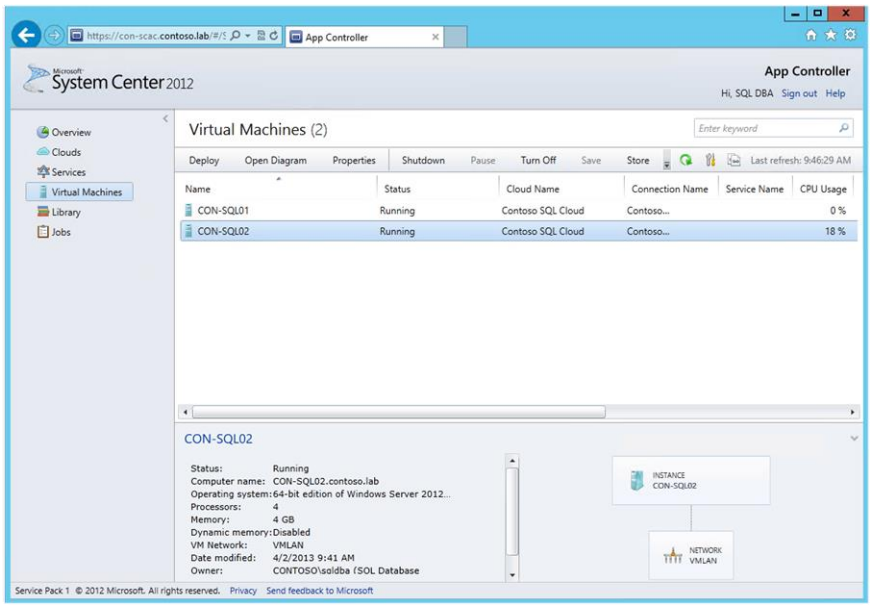

By selecting a particular cloud, the database administrator can choose from a list of service templates and, from there, provide the final pieces of configuration to customize a deployment, such as Service Name,

Virtual Machine Name, and Operating System Name. The DBA clicks *Deploy* to start the virtual machine provisioning process, and VMM automatically orchestrates the deployment and placement of the SQL Server virtual machine (Figure 56).

Figure 56: Deploying a SQL Server virtual machine with App Controller

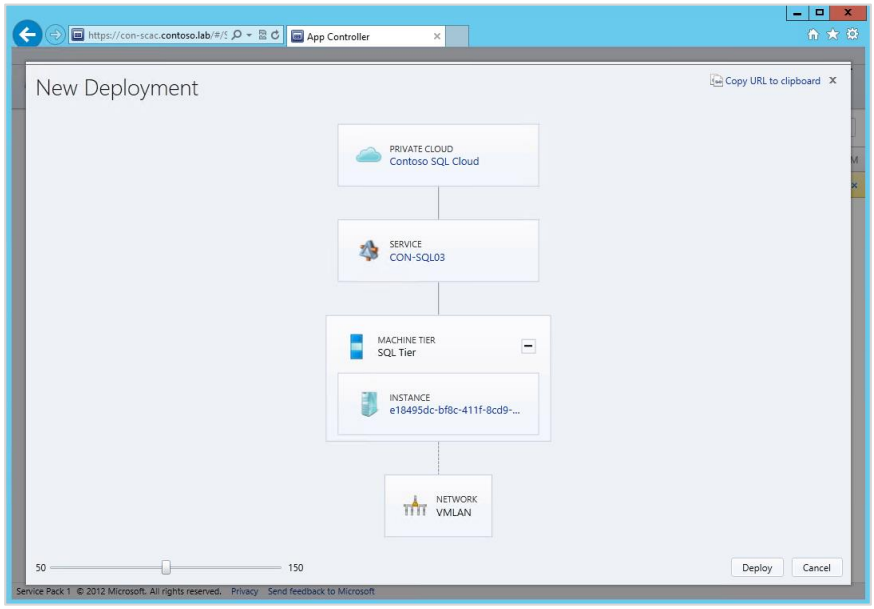

Once the new virtual machine is deployed, the database administrator can access it through App Controller and perform the tasks and actions that the IT administrators have enabled. In the example in Figure 57, this includes the ability for the DBA to connect to the virtual machine via remote desktop to perform any SQL-specific actions.

Figure 57: Connecting to a virtual machine through App Controller

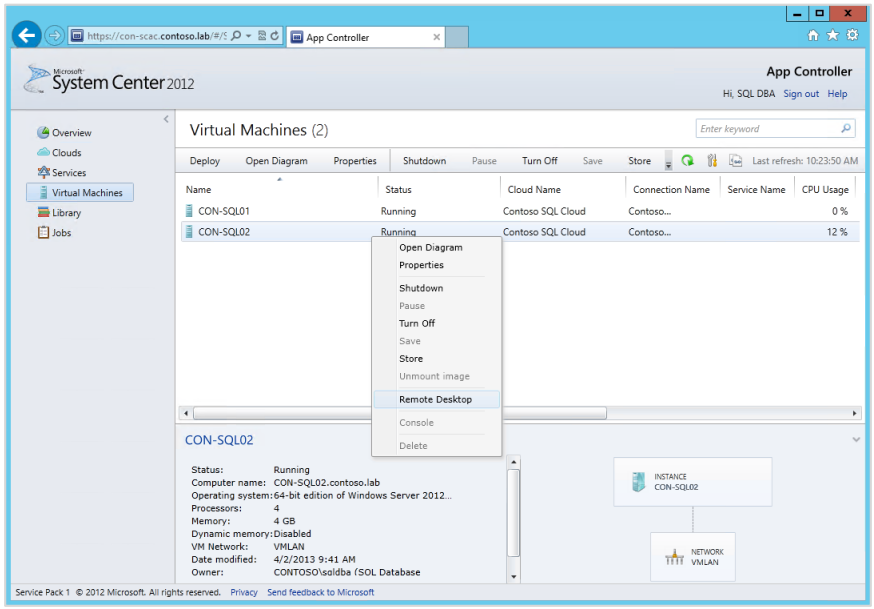

Service Delivery and Automation

To review, the subsections above have discussed:

- How IT administrators can define a private cloud within VMM, generate templates, and assign users/groups.
- How, from that point forward, database administrators can access the rich web interface of App Controller to deploy virtual machines and services into that cloud.

Now consider the scenario where IT administrators want to use App Controller for DBAs to access their virtual machines, but also want to enact a mechanism through which the DBAs must request new virtual machines, as needed, instead of creating them at will. To manage this scenario, your organization needs the Service Manager and Orchestrator components of System Center, as well as the free Cloud Services Process Pack download. Together, these elements, along with other System Center components like VMM, deliver a self-service infrastructure-as-a-service (IaaS) platform that is managed by IT and consumed by end users, application owners, and DBAs.

Before examining how the components work together, it is important to understand what they provide individually. Each component is discussed below in more detail.

Service Manager

- **IT Service Management**: System Center 2012 SP1 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.
- **ITaaS**: Service Manager enables a rich self-service portal that provides role-based access to the service catalog (Figure 58). The Self-Service Portal in System Center 2012 is a SharePoint website that is accompanied by a set of Microsoft Silverlight applications. The SharePoint environment provides a foundation on which the portal can be customized. It also provides a set of building blocks for extending the features that users can access through a web browser.

Figure 58: Self-service portal in Service Manager

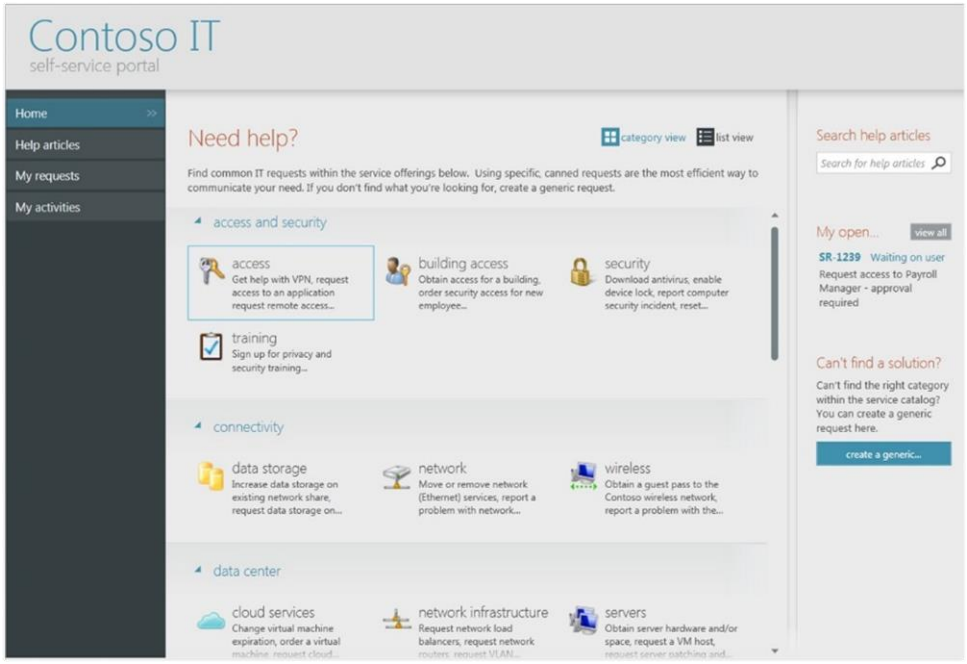

- **Integration**: Connectors simplify and streamline integration between Service Manager and other System Center components. You can use Service Manager connectors to import data as configuration items from Active Directory Domain Services, Configuration Manager, Orchestrator, VMM, and Operations Manager. In addition, you can import alerts from Operations Manager and configure them to automatically generate incidents in Service Manager. You can also import data from comma-separated value (CSV) files into the Service Manager database.
- **Business Intelligence**: Service Manager delivers a powerful data warehouse for rich, integrated reporting. Service Manager reports enable you to collect and view data and trends from across the business environment. For example, you can generate a report that shows the number of incidents that occur in a specific time frame. You can then use that information to calculate the cost of each incident (in hours) and to identify trends and take preventative measures to reduce the cost and occurrence of incidences.

Orchestrator

- **Custom Automation**: System Center 2012 SP1 Orchestrator provides tools to build, test, debug, deploy, and manage automation in your environment. These automated procedures, called *runbooks*, can function independently or start other runbooks (Figure 59). The standard activities defined in every installation of Orchestrator provide a variety of monitors, tasks, and runbook controls, which you can integrate with a wide range of system processes. Each activity in a runbook publishes data that is available to any subsequent activity in that runbook. You can use this published data to provide dynamic decision-making capabilities (like creating emails, alerts, log files, accounts, and more).

Figure 59: Sample runbook in Orchestrator

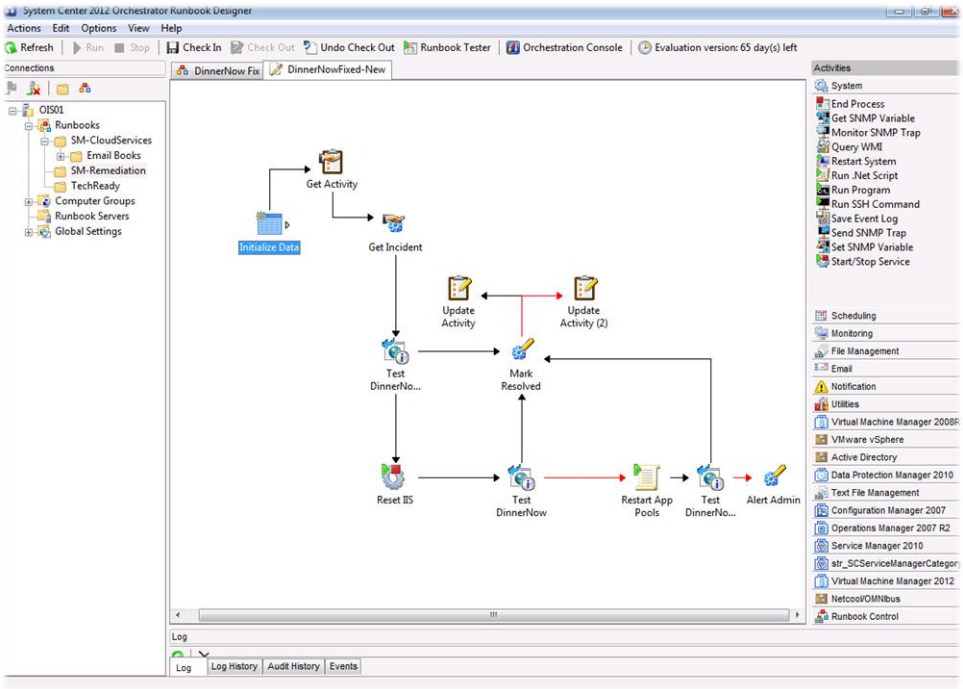

Your IT organization can use Orchestrator to improve efficiency and reduce operational costs to support cross-departmental objectives. Orchestrator provides an environment with shared access to common data. By using Orchestrator, you can evolve and automate key processes between groups and consolidate repetitive manual tasks. You can automate cross-functional team processes and enforce best practices for incident, change, and service management by creating runbooks that are customized for your requirements. Through automation, regularly recurring tasks reduce the number of manual and error-prone activities in your environment, helping to improve reliability and predictability.

- **Cross-Platform Integration**: Orchestrator integrates with System Center, other Microsoft products, and non-Microsoft products to enable interoperability across the data center. Orchestrator improves efficiency across multiple tools, systems, and departments by eliminating or crossing technology and organizational process structures. You can extend the capabilities of Orchestrator with integration packs that include additional functionality for both Microsoft and non-Microsoft products and technologies. Orchestrator activities and integration packs reduce unanticipated errors and shorten service delivery time by automating the common tasks associated with enterprise tools and products.
- **End-to-End Orchestration**: *Orchestration* is the collective name for the automated arrangement, coordination, and management of systems, software, and practices. It enables the management of complex cross-domain processes. Orchestrator provides the tools for orchestration to combine software, hardware, and manual processes into a seamless system. These tools let you connect and automate workflows.

  Just as manufacturing companies have automated common and repeatable tasks from their production processes, you can adopt this same efficiency in the IT environment by using Orchestrator to seamlessly perform and monitor your IT processes. Orchestrator can handle routine tasks, ensure process enforcement, and reliably meet the demands of the largest

enterprises. Orchestrator interoperates with other System Center products to integrate IT administrative tasks from start to finish.

- **Extensible Structure**: If you have a custom in-house solution, Orchestrator provides extensible integration to any system through the Orchestrator Integration Toolkit. You can create custom integrations that allow Orchestrator to connect to any environment. Orchestrator uses a Representational State Transfer (REST)-based web service that can perform processes like start and stop runbook jobs and get reporting information in Open Data protocol (OData) format. The web service lets you develop applications that can use live data from Orchestrator.

Cloud Services Process Pack

- **Infrastructure as a Service**: IaaS is a service-centric model for requesting and provisioning data center resources. The System Center Cloud Services Process Pack is the Microsoft IaaS solution built on the System Center platform. With the Cloud Services Process Pack, your organization can realize the benefits of IaaS while simultaneously using your existing investments in Service Manager, Orchestrator, VMM, and Operations Manager.

  Corporate data centers are in transition. The recent shift from physical to virtual environments is now being replaced by an interest in moving to the cloud—specifically both private and public cloud infrastructures. Private cloud management assets are being delivered with Service Manager, and a key part of this solution is the self-service experience. This experience is now significantly enhanced by the Cloud Services Process Pack.

  Moreover, IT organizations considering IaaS need to examine and adapt their existing tools, processes, workflows, and automation to meet the requirements of an effective cloud services implementation. While it is critical that the underlying features (such as the Self-Service Portal, ticketing infrastructure, notifications, workflows, and automation) integrate well with each other and account for industry-wide best practices, the work involved to ensure an effective cloud services implementation can be daunting and time-consuming. The Cloud Services Process Pack addresses these concerns by enabling IaaS while incorporating domain expertise and best practices from organizations that have successfully deployed IaaS.

The IaaS Solution

The Service Manager, Orchestrator, and Cloud Services Process Pack components work together to form a powerful IaaS solution. With this solution, designated users like database administrators can request infrastructure; once the request is approved, integrated automation orchestrates delivery of and access to the infrastructure—reducing the need for IT involvement and accelerating time to market.

Using the key components of System Center and the Cloud Services Process Pack, IT administrators can define a rich self-service experience for DBAs who want to request infrastructure to run their database workloads. In Figure 60, when the DBA logs on to the Contoso Portal, the portal recognizes who the user is. Role-based access is key to the Service Manager self-service experience, and the portal dynamically generates content based on the specific user.

Figure 60: Self-service portal in Service Manager

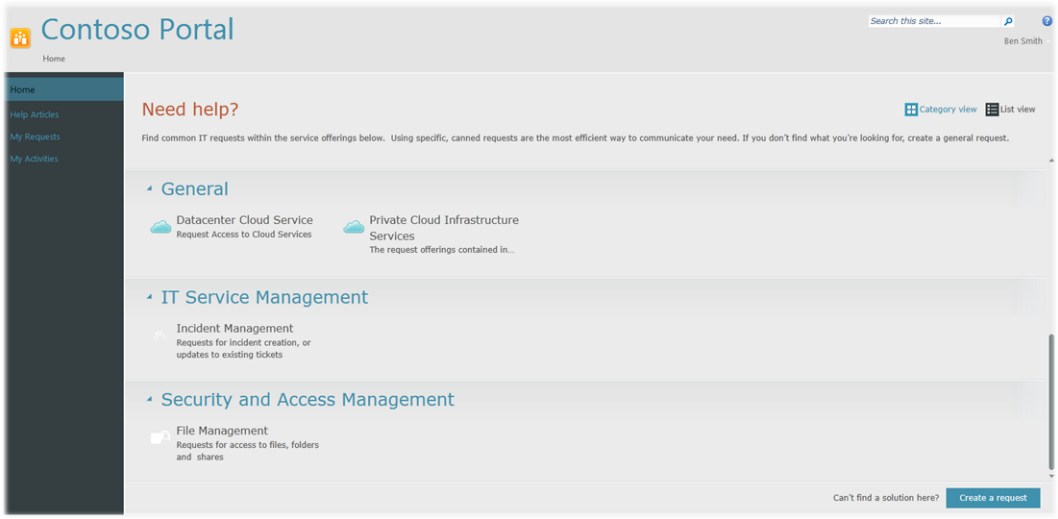

Figure 61 shows the Service Offerings page of the portal. Service Offerings essentially group together requests that the specific user can make. In this example, the DBA selects the Service Offering entitled *Private Cloud Infrastructure Services* to be presented with available requests.

Figure 61: Service Offerings page and related requests in Service Manager

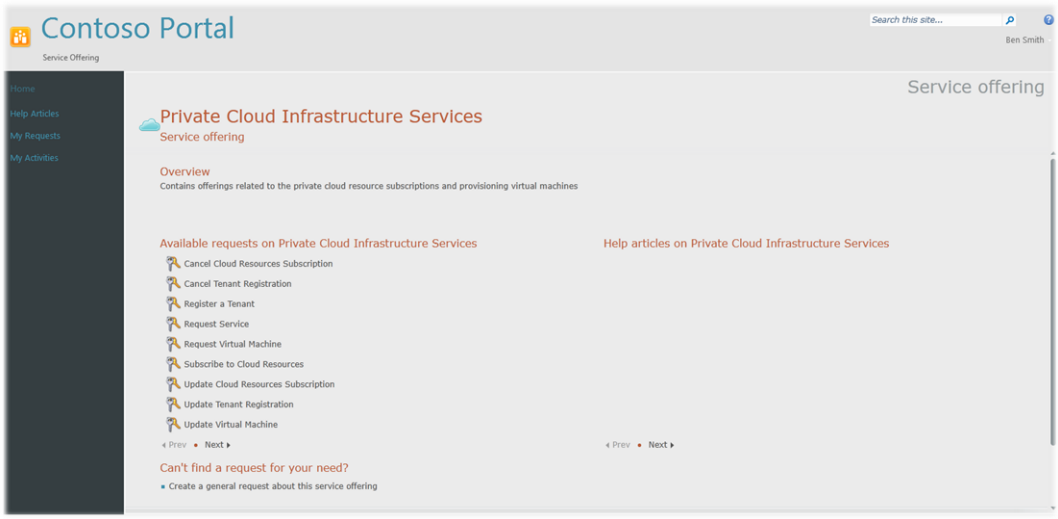

The available requests are essentially the menu of choices that IT has provided for the DBA. The Cloud Services Process Pack provides all of these in the box, and they can be used as templates for further customization by IT. Example requests include:

- Tenant Registration
- Tenant Update Registration
- Cloud Resources Subscription
- Cloud Resources Update Subscription
- Virtual Machine

- Virtual Machine Update
- Tenant Registration Cancellation
- Cloud Resources Subscription Cancellation

From a SQL perspective, IT can define specific requests that relate to SQL Server. These requests can be general or specific—for example, a request for IT to create a pool of resources for the SQL database team, or a request to deploy a specific SQL virtual machine. Remember the previous example where the SQL DBA used App Controller to deploy a virtual machine from a service template. In the current example, the same process is taking place, but the DBA is making a request to have the activity performed (Figure 62). If the request is approved, Orchestrator works in conjunction with Service Manager and VMM to create and deliver the virtual machine. From there, the DBA can interact with the virtual machine through App Controller.

Figure 62: Making a cloud resource request in Service Manager

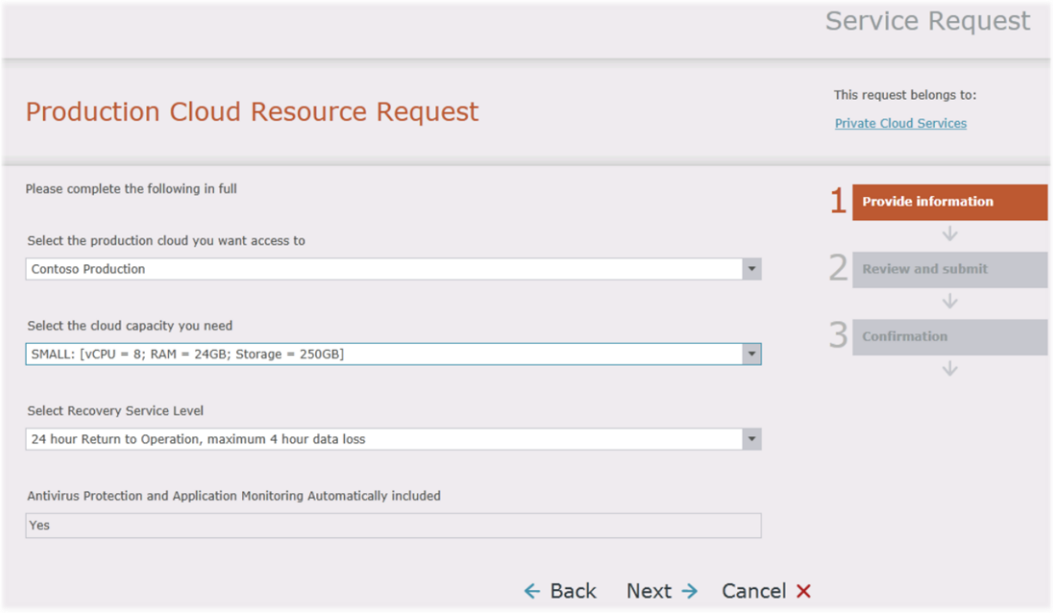

Note that the DBA can request a private cloud and specify its capacity, recovery service level, and antivirus and monitoring status. Once the request is submitted, the appropriate IT administrator is notified to initiate approval. Upon approval, Orchestrator, as part of the integrated CSPP runbooks, plugs in the relevant information from the form; orchestrates the creation of the cloud with VMM; and sets up the relevant monitoring and protection with Operations Manager and Data Protection Manager. The DBA receives notification upon completion and then is able to access the resources with App Controller or through remote desktop.

Operations Manager

Microsoft has a long history of defining and refining monitoring capabilities for its products. System Center 2012 SP1 Operations Manager continues this tradition as a solution with a deeper level of insight and improved scalability. With Operations Manager, your organization can gain levels of visibility into its infrastructure at every level of the stack, helping to ensure that the infrastructure is optimized and running efficiently. Fundamentally, Operations Manager provides infrastructure monitoring that is flexible and cost effective, better ensures the predictable performance and availability of vital applications, and offers comprehensive oversight of your data center and cloud—both private and public.

Operations Manager enables IT administrators to monitor services, devices, and operations for many computers in a single console (Figure 63). Operations Manager includes numerous views that show state, health, and performance information, as well as alerts generated for availability, performance, configuration, and security situations. With these tools, you can gain rapid insight into the state of the IT environment and the IT services running across different systems and workloads.

Figure 63: Dashboard in Operations Manager

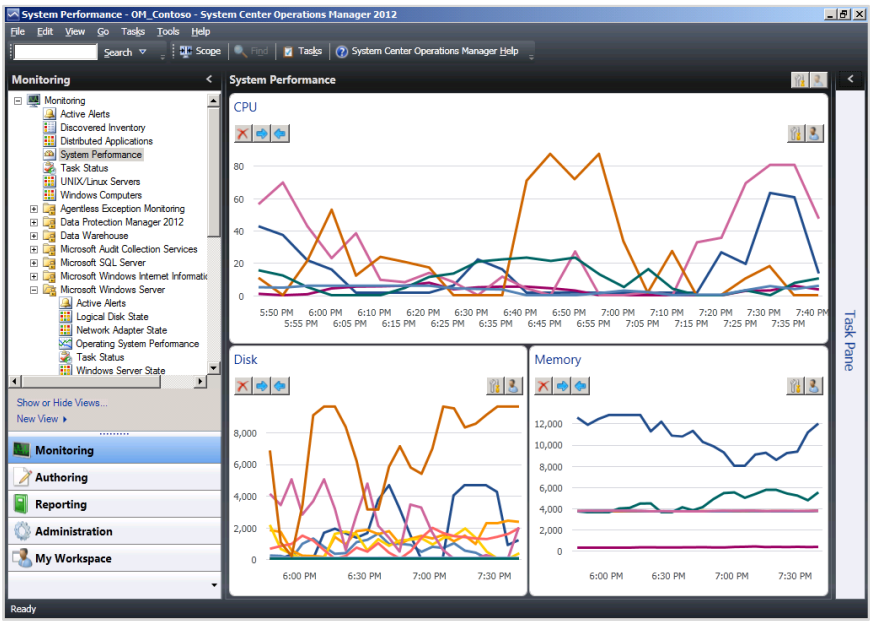

SQL Server 2012 Management Pack

The ability to provide end-to-end management across infrastructure is a critical step in ensuring the health of the hosts and clusters, virtual machines, and private cloud itself. Further, Operations Manager can be used to monitor mission-critical SQL Server workloads using the SQL Server 2012 Management Pack. This management pack provides the capabilities for Operations Manager 2007 R2 and Operations Manager 2012 to discover SQL Server 2005, 2008, 2008 R2, and SQL Server 2012. It monitors SQL Server components such as database engine instances, databases, and SQL Server agents.

The SQL Server 2012 Management Pack monitors performance, availability, and configuration; performance data collection; and default thresholds. In addition to health monitoring capabilities, this management pack includes dashboard views, diagram views and extensive knowledge with embedded

inline tasks, and views that enable near real-time diagnosis and resolution of detected issues. Plus, you can integrate the monitoring of SQL Server components into your service-oriented monitoring scenarios.

The SQL Server 2012 Management Pack provides the following key features:

- **SQL Version Support**: Support for Enterprise, Standard, and Express editions of SQL Server 2005, 2008, 2008 R2, and 2012, as well as 32bit, 64bit, and ia64 architectures.
- **Simple and Complex Configuration Support**: Clustered installations, multiple instances, and 32bit roles running on a 64bit operating system.
- **Discovery and Monitoring**: SQL Server roles such as DB Engine, Reporting Services, Analysis Services, and Integrations Services, as well as SQL Server components such as databases, SQL Agent, and SQL jobs.
- **Core Monitoring**: Database free space, SQL Server-related performance, SQL Server-related alerts, and lists of the various SQL Server roles and components that are discovered and their related states.
- **AlwaysOn Monitoring**: Automatic discovery and monitoring of availability groups, availability replicas, and availability databases for hundreds of computers; health roll-up from availability database to availability replicas; and detailed knowledge with every critical health state to enable faster resolution to a problem.
- **Policy-Based Management**: Auto-discovery of custom PBM polices targeting AlwaysOn and database components; roll-up of the health of policy execution within the management pack under extended health.
- **Mirroring and Replication Monitoring**: Discovery of mirroring databases, witness, and mirroring group; monitoring of database mirror state, database mirror witness state, and mirroring partners' state; custom diagram view to visually represent the primary and mirrored databases; approximately 20 rules to detect replication events.

By deploying Operations Manager with the SQL Server 2012 Management Pack, IT administrators and database administrators can gain a deeper level of insight into their workloads, helping to optimize performance and remediate issues quickly. Continuing the DBA example, a DBA can access this information directly in the Operations Manager console using the *My Workspace* functionality. My Workspace provides a private area in the Operations Manager console that can be customized for the DBA's specific needs. Using My Workspace, the DBA can create folders to organize the workspace, add shortcuts to favorite views, save useful searches, and create views that are only visible to self. Using the same Windows credentials, the DBA can log on to My Workspace from any Operations Manager console. In addition, the DBA can subscribe to alerts from Operations Manager to receive email notification of any issues that arise.

Best Practices and Recommendations

If the DBA does not want to install the full Operations Manager console on a local machine or if IT is not comfortable providing that level of access to the DBA, the Operations Manager web console is a good option (Figure 64). The web console can be accessed through a browser and offers near-parity with the full console. This allows the DBA to access key information, alerts, dashboards, and views focused on specific areas (such as SQL Server).

Figure 64: Web console for Operations Manager

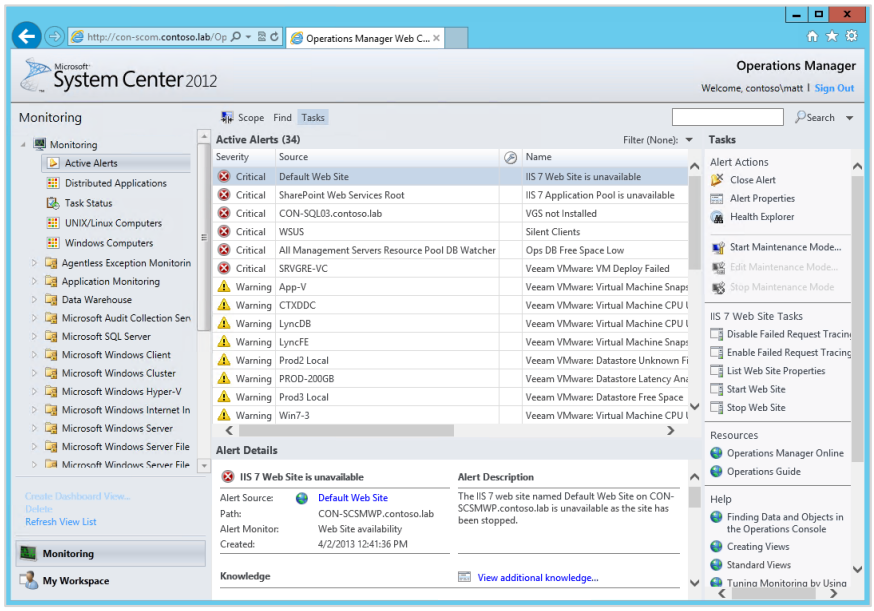

Data Protection Manager

Data Protection Manager (DPM) has been part of the System Center family since the 2006 version was launched, and since then, it has continued to evolve with more enhanced capabilities around backup and workload protection. With the 2012 SP1 release, there have been further incremental improvements in DPM that make it an excellent companion for Hyper-V, offering granular, efficient protection of virtual machines (Figure 65).

Figure 65: Key capabilities of Data Protection Manager

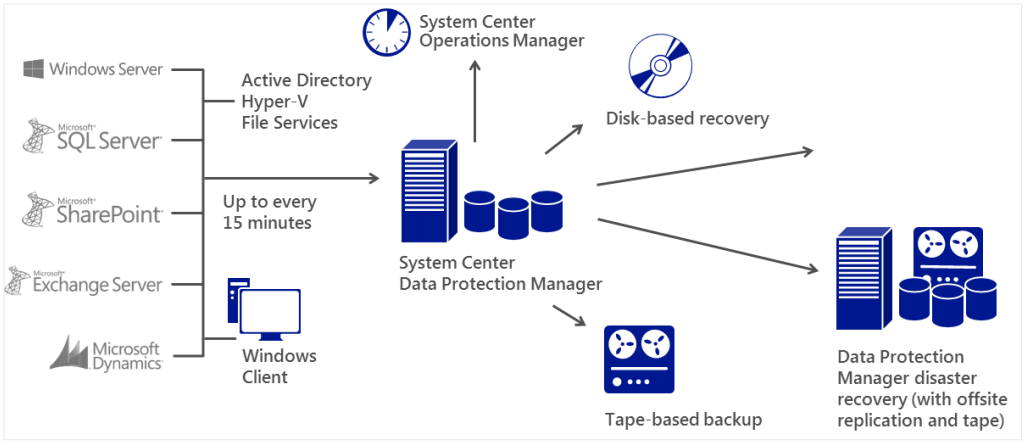

DPM provides continuous data protection for SQL Server, protecting both standalone SQL instance databases and those that use high availability technologies such as transaction log replication, clustering support, database mirroring, and AlwaysOn Availability Groups in SQL Server 2012. Likewise, DPM offers comprehensive SQL Server data protection for organizations of all sizes through disk-to-disk, disk-to-tape, and disk-to-disk-to-tape technologies, helping to maintain the business value of a SQL infrastructure by ensuring that it is better protected and always available.

With System Center 2012 SP1, DPM now can back up data from the DPM server to offsite storage that is managed by the Windows Azure Online Backup service. (Your organization must sign up for the service, and you must download and install the Windows Azure Online Backup agent on the DPM server, which is used to transfer the data between the server and the service.) With the new online backup capabilities in DPM, you can expect the following benefits:

- **Reduced TCO**: The Windows Azure Online Backup service can help to reduce TCO by providing scalability, elasticity, and simplified storage management.
- **Peace of Mind**: The Windows Azure Online Backup service provides a reliable, secure, and robust offsite solution for backup and recovery that is highly available.
- **Simplicity**: The Windows Azure Online Backup workflows are seamlessly integrated into the existing DPM backup, recovery, and monitoring workflows.

Non-Microsoft backup solutions tend to take generic backup functionality and adapt it to support specific applications. In contrast, Microsoft created DPM to capitalize on fully supported Microsoft technologies in order to provide near-continuous data protection. This support for SQL Server includes:

- SQL instance-level protection.
- Data source co-location, which allows a DPM server to protect more than 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 deployment without interrupting access to databases:

- Baseline copy of the SQL Server data can be made using either the DPM block-level synchronization engine or manual intervention.
- 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. (Note that a best practice is to configure express full backups every evening and to synchronize transaction logs more frequently—every 15 to 60 minutes.)

In terms of recovery, DPM supports the following SQL Server recovery scenarios:

- Recovering all protected SQL Server data.
- Recovering a specific database to its original location.
- Recovering a specific database to an alternate instance.
- Recovering database files to a network folder or their own tape.

These options provide administrators with flexible choices around restoring databases in case of failure or corruption. For delegated DBAs, IT administrators can define self-service recovery capabilities for restoring SQL Server databases in the event of loss.

Best Practices and Recommendations

DPM has the following considerations for SQL Server 2012 protection:[53]

- Users must explicitly add the system account *NTAuthority\System* to the Sysadmin group on SQL Server.
- Availability Group replicas must be configured as read-only.
- When an alternate location recovery is performed for a partially contained database, the target SQL instance must have the Contained Databases feature enabled.
- DPM does not support protection for SQL Server data on a remote file share. If there is a database with files on a remote file share, protection will fail with Error ID 104.
- DPM does not honor the backup policy that is set in SQL Server.
- DPM will never back up from the primary server. If there is no secondary server, DPM will not show the primary server. If there are three replicas (one primary server and two secondary servers), DPM will show only the secondary servers.
- DPM cannot back up from an asynchronous secondary server.
- Recovery to original location is not supported.

Conclusion

Virtualizing SQL Server 2012 on Windows Server 2012 Hyper-V helps to ensure optimal performance and ease of management. Proper planning is required before virtualizing any mission-critical workload like SQL Server, and it is beneficial to understand the best practices and recommendations detailed in this guide. At a high level, the fabric considerations can help you to effectively plan the physical infrastructure, including processors, memory, storage, and network. Likewise, the resiliency considerations can help you negotiate the configuration of virtual machines, SQL Server 2012, and Windows Server 2012 Hyper-V according to various requirements.

By choosing Windows Server 2012 Hyper-V to virtualize SQL Server 2012, your organization can significantly lower TCO and operational costs, improve hardware capabilities, and maximize ROI. Further, your organization can consolidate SQL Server 2012 workloads using virtualization to address issues related to server sprawl; increase ROI on database workloads; and control costs (including operations and energy). System Center 2012 SP1 can be used to proactively manage, monitor, and secure the SQL Server virtualization environment.

Together, SQL Server 2012, Windows Server 2012, and System Center 2012 SP1 provide an integrated information platform spanning both on-premises data centers and the cloud. This platform provides a powerful approach and customizable options to meeting the complex and growing demands of today's enterprises.

Additional Resources

For more information, please visit the following links:

Microsoft SQL Server 2012
http://www.microsoft.com/en-us/sqlserver/default.aspx

Microsoft SQL Server 2012 TechNet
http://technet.microsoft.com/en-us/sqlserver/ff898410.aspx

Windows Server 2012
http://www.microsoft.com/en-us/server-cloud/windows-server/default.aspx

Windows Server 2012 TechNet
http://technet.microsoft.com/en-us/windowsserver/hh534429.aspx

Microsoft System Center 2012
http://www.microsoft.com/en-us/server-cloud/system-center/default.aspx

Microsoft System Center 2012 TechNet
http://technet.microsoft.com/en-us/systemcenter/bb980621.aspx

What's New in System Center 2012 SP1
http://technet.microsoft.com/en-us/systemcenter/bb980621.aspx

References

[1] Microsoft. "High Performance SQL Server Workloads on Hyper-V." May 2010.
http://download.microsoft.com/download/d/f/8/df89d22d-39c8-4728-a990-3bd4467891b7/highperformancesqlserverworkloadsonHyper-V_final.docx

[2] Leon, Mike. "Microsoft ESG Lab Validation Report." Oct 2012.
http://download.microsoft.com/download/C/C/1/CC16C89A-E289-4217-B2D8-7DD37A4285B8/ESG-Lab-Validation-WS2012-HyperV-and-SQL2012.pdf

[3] Microsoft. "Competitive Advantages of Windows Server 2012 Hyper-V over VMware vSphere 5.1."
http://download.microsoft.com/download/E/8/E/E8ECBD78-F07A-4A6F-9401-AA1760ED6985/Competitive-Advantages-of-Windows-Server-Hyper-V-over-VMware-vSphere.pdf

[4] Microsoft TechNet Blog. "The Perfect Combination: SQL Server 2012, Windows Server 2012 and System Center 2012." Dec 2012.
http://blogs.technet.com/b/dataplatforminsider/archive/2012/12/06/the-perfect-combination-sql-server-2012-windows-server-2012-and-system-center-2012.aspx

[5] Microsoft. "System Center 2012 - Infrastructure Management."
http://www.microsoft.com/en-us/server-cloud/system-center/infrastructure-management.aspx

[6] Microsoft. "Installing Windows Server 2012." May 2012.
http://technet.microsoft.com/en-us/library/jj134246.aspx

[7] Microsoft. "High Performance SQL Server Workloads on Hyper-V." May 2010.
http://download.microsoft.com/download/d/f/8/df89d22d-39c8-4728-a990-3bd4467891b7/highperformancesqlserverworkloadsonHyper-V_final.docx

[8] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[9] Coombes, David. "Module 1 - VM Scale Student Manual." Sep 2012
http://download.microsoft.com/download/F/F/1/FF10F038-26CD-4B00-858C-C99C0D9FAB93/Module%201%20-%20VM%20Scale%20Student%20Manual.pdf

[10] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[11] Microsoft TechNet Blog. "Hyper-V VM Density, VP:LP Ratio, Cores and Threads." Apr 2011.
http://blogs.technet.com/b/virtualization/archive/2011/04/25/Hyper-V-vm-density-vp-lp-ratio-cores-and-threads.aspx

[12] Armstrong, Ben (MSDN Blog). "Hyper-V CPU Scheduling–Part 1." Feb 2011.
http://blogs.msdn.com/b/virtual_pc_guy/archive/2011/02/14/hyper-v-cpu-scheduling-part-1.aspx

[13] Microsoft MSDN. "Understanding Non-Uniform Memory Access."
http://msdn.microsoft.com/en-us/library/ms178144.aspx

[14] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[15] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[16] Microsoft MSDN. "Performance Tuning Guidelines for Windows Server 2012." Oct 2012.
http://download.microsoft.com/download/0/0/B/00BE76AF-D340-4759-8ECD-C80BC53B6231/performance-tuning-guidelines-windows-server-2012.docx

[17] ten Seldam, Matthijs (TechNet Blog). "Windows Server - Sockets, Logical Processors, Symmetric Multi-Threading." Oct 2012.
http://blogs.technet.com/b/matthts/archive/2012/10/14/windows-server-sockets-logical-processors-symmetric-multi-threading.aspx

[18] Microsoft TechNet. "Storage Spaces - Designing for Performance."
http://social.technet.microsoft.com/wiki/contents/articles/15200.storage-spaces-designing-for-performance.aspx

[19] Microsoft TechNet. "Storage Spaces - Designing for Performance."
http://social.technet.microsoft.com/wiki/contents/articles/15200.storage-spaces-designing-for-performance.aspx

[20] Microsoft. "Microsoft ESG Lab Validation - Windows Server Storage." Dec 2012.
http://download.microsoft.com/download/8/0/F/80FCCBEF-BC4D-4B84-950B-07FBE31022B4/ESG-Lab-Validation-Windows-Server-Storage.pdf

[21] Barreto, Jose (TechNet Blog). "Hyper-V over SMB - Sample Configurations."
http://blogs.technet.com/b/josebda/archive/2013/01/26/hyper-v-over-smb-sample-configurations.aspx

[22] Microsoft TechNet. "Failover Clustering Hardware Requirements and Storage Options." Aug 2012.
http://technet.microsoft.com/en-in/library/jj612869.aspx

[23] Microsoft. "WS 2012 Networking Whitepaper." http://download.microsoft.com/download/7/E/6/7E63DE77-EBA9-4F2E-81D3-9FC328CD93C4/WS%202012%20White%20Paper_Networking.pdf

[24] Microsoft TechNet. "Building Your Cloud Infrastructure: Converged Data Center with File Server Storage. Sep 2012.
http://technet.microsoft.com/en-us/library/hh831738.aspx

[25] Microsoft. "Windows Server 2012 Product Overview."
http://download.microsoft.com/download/B/3/0/B301FE04-ABAB-4C89-8631-0FCBE4147B84/WS_2012_Product_White_Paper.pdf

[26] Microsoft TechNet. "Cluster-Aware Updating Overview."
http://technet.microsoft.com/en-us/library/hh831694.aspx

[27] Microsoft TechNet. "Building Your Cloud Infrastructure: Non-Converged Data Center Configuration."
http://technet.microsoft.com/en-us/library/hh831559.aspx

[28] Microsoft TechNet. "Configuring Availability Options for Virtual Machines in System Center 2012 SP1 - Overview." Jan 2013.
http://technet.microsoft.com/en-us/library/jj628163.aspx

[29] Microsoft TechNet. "What's New in Failover Clustering." Sep 2012.
http://technet.microsoft.com/en-us/library/hh831414.aspx#BKMK_PLACE

[30] Microsoft. "Support Policy for SQL Server Products that are Running in a Hardware Virtualization Environment. " Feb 2013.
http://support.microsoft.com/kb/956893

[31] Hsueh, Sung, et al. "SQL Server Consolidation Guidance." Nov 2009.
http://msdn.microsoft.com/en-us/library/ee819082.aspx

[32] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[33] Microsoft MSDN. "Performance Tuning Guidelines for Windows Server 2012." Oct 2012.
http://download.microsoft.com/download/0/0/B/00BE76AF-D340-4759-8ECD-C80BC53B6231/performance-tuning-guidelines-windows-server-2012.docx

[34] Allen, Lindsey, et al. "Running SQL Server 2008 in a Hyper-V Environment." Oct 2008.
http://download.microsoft.com/download/d/9/4/d948f981-926e-40fa-a026-5bfcf076d9b9/sql2008inhyperv2008.docx

[35] Posey, Brien (TechNet Blog). "Virtualization: Optimizing Hyper-V Memory Usage."
http://technet.microsoft.com/en-us/magazine/hh750394.aspx

[36] Microsoft. "Implementing and Configuring Dynamic Memory." June 2010.
http://download.microsoft.com/download/E/0/5/E05DF049-8220-4AEE-818B-786ADD9B434E/Implementing_and_Configuring_Dynamic_Memory.docx

[37] Microsoft TechNet. "Hyper-V Dynamic Memory Overview." Feb 2012.
http://technet.microsoft.com/en-us/library/hh831766.aspx

[38] Coombes, David. "Module 1 - VM Scale Student Manual." Sep 2012.
http://download.microsoft.com/download/F/F/1/FF10F038-26CD-4B00-858C-C99C0D9FAB93/Module%201%20-%20VM%20Scale%20Student%20Manual.pdf

[39] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[40] Microsoft. "WS2012 Whitepaper Hyper-V."
http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

[41] Microsoft TechNet. "Configure the Memory for the Virtual Machines." Jan 2013
http://technet.microsoft.com/en-us/library/ff621103.aspx#ConfigMem

[42] Microsoft. "Running SQL Server with Hyper-V Dynamic Memory." Jul 2011
http://download.microsoft.com/download/D/2/0/D20E1C5F-72EA-4505-9F26-FEF9550EFD44/Best%20Practices%20for%20Running%20SQL%20Server%20with%20HVDM.docx

[43] Rabeler, Carl. "Running Microsoft SQL Server 2008 Analysis Services on Windows Server 2008 vs. Windows Server 2003 and Memory Preallocation: Lessons Learned." Jul 2008.
http://sqlcat.com/sqlcat/b/technicalnotes/archive/2008/07/16/running-microsoft-sql-server-2008-analysis-services-on-windows-server-2008-vs-windows-server-2003-and-memory-preallocation-lessons-learned.aspx

[44] Microsoft Support. "Hard Disk Drive Sector-Size Support Boundaries in SQL Server."
http://support.microsoft.com/kb/926930

[45] Microsoft TechNet. "Optimizing Performance on Hyper-V."
http://technet.microsoft.com/en-us/library/cc768529(v=bts.10).aspx

[46] Microsoft TechNet. "Everything You Wanted to Know about SR-IOV in Hyper-V (Part 5)." Mar 2012.
http://blogs.technet.com/b/jhoward/archive/2012/03/16/everything-you-wanted-to-know-about-sr-iov-in-Hyper-V-part-5.aspx

[47] Microsoft TechNet. "Quality of Service (QoS) Overview." Mar 2013.
http://technet.microsoft.com/en-us/library/hh831679.aspx#bkmk_bandwidth

[48] Microsoft TechNet. "QoS Minimum Bandwidth Best Practices." Oct 2012.
http://technet.microsoft.com/en-us/library/jj735303.aspx

[49] Microsoft TechNet. "AlwaysOn Failover Cluster Instances (SQL Server)."
http://technet.microsoft.com/en-us/library/ms189134.aspx

[50] Microsoft TechNet. "The Perfect Combination: SQL Server 2012, Windows Server 2012 and System Center 2012."
http://blogs.technet.com/b/dataplatforminsider/archive/2012/12/06/the-perfect-combination-sql-server-2012-windows-server-2012-and-system-center-2012.aspx

[51] Microsoft MSDN. "Windows Server Failover Clustering (WSFC) with SQL Server."
http://msdn.microsoft.com/en-us/library/hh270278.aspx

[52] Microsoft MSDN. "Overview of AlwaysOn Availability Groups (SQL Server)."
http://msdn.microsoft.com/en-us/library/ff877884.aspx

[53] Microsoft TechNet. "Protecting SQL Server Data." Jan 2013.
http://technet.microsoft.com/en-us/library/hh780998.aspx