



# Best Practices for Virtualizing & Managing SharePoint 2013

v1.0 – May 2013



# Copyright Information

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

# Table of Contents

Copyright Information .....	2
Table of Contents .....	3
Introduction .....	5
Executive Summary .....	5
Target Audience .....	6
Scope .....	6
Why Virtualize SharePoint? .....	7
Why Microsoft Virtualization & Management? .....	9
Fabric Configuration .....	10
Hardware Considerations .....	10
Compute Considerations .....	12
Storage Considerations .....	18
Networking Considerations .....	30
Host Resiliency & VM Agility .....	36
Host Clustering .....	36
Virtual Machine Configuration .....	45
Microsoft Assessment and Planning Toolkit .....	45
SharePoint Virtual Machine CPU Considerations .....	45
SharePoint Virtual Machine Memory Considerations .....	47
SharePoint Virtual Machine Storage Considerations .....	49
SharePoint Virtual Machine Networking Considerations .....	52
SharePoint Virtual Topologies .....	59

System Center 2012 SP1 .....	66
Comprehensive Management Capabilities.....	66
Virtual Machine Manager .....	66
App Controller .....	75
Service Delivery and Automation .....	77
Operations Manager .....	83
Data Protection Manager .....	84
Conclusion .....	86
Additional Resources.....	87
References.....	88

# Introduction

This guide provides high-level best practices and considerations for deploying and managing Microsoft SharePoint 2013 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 SharePoint 2013 to deliver the highest levels of performance and availability.

## Executive Summary

As information grows within organizations, it becomes increasingly difficult to effectively organize, manage, and share. Microsoft SharePoint has become a critical resource for many organizations, providing new ways to share and collaborate, organize projects and team, and discover people and information.

To meet the demands of ever-changing business needs and fast-paced collaboration, SharePoint must be rapidly scalable and highly available at all times. As demand for SharePoint increases within and across organizations, IT must quickly meet these requirements, while also minimizing the cost of the underlying infrastructure, managing risks, and saving time and effort on daily management tasks. An ideal solution to these challenges is virtualization.

Virtualization is increasingly common in today's corporate environments. Organizations worldwide have moved beyond the early stage of virtualization into its more advanced uses. Server virtualization is a specific area that can provide robust benefits. Organizations that have virtualized their servers have seen notable improvements in cost, efficiency, operations, availability, agility, and resiliency.

SharePoint 2013 can help organizations manage cost and gain efficiency by providing features that enhance the flexibility and scalability of SharePoint infrastructure. Likewise, with its enhanced capabilities, Windows Server 2012 Hyper-V can be used to effectively virtualize mission-critical workloads, such as SharePoint 2013. This guide explores key product features and capabilities in the context of virtualizing SharePoint 2013 on Windows Server 2012. It also examines the benefits of managing this virtualized environment with Microsoft System Center 2012 SP1. Working together, these powerful products deliver an integrated platform that offers low total cost of ownership (TCO), mission-critical scale and performance, and high availability. The platform also provides more 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 SharePoint 2013.

# Target Audience

This guide is intended for IT professionals and technical decision makers (TDMs), including IT architects, managers, and administrators. They can use this guide to understand how to set up an environment for virtualizing SharePoint 2013 using an integrated virtualization platform built on some of the latest Microsoft technologies, including Windows Server 2012 Hyper-V and System Center 2012 SP1. Understanding key considerations and best practices can help TDMs effectively plan and deploy SharePoint 2013 in a virtualized environment. This guide serves the following purposes for these key IT roles:

- **Architects:** Understand how the entire virtualization environment will work as they design the architecture and how specific underlying features can ensure the highest levels of performance of the SharePoint 2013 infrastructure.
- **Managers:** Design processes to fit the overall virtualization environment so that costs are reduced and efficiency is increased as much as possible.
- **Administrators:** Understand important considerations while managing the virtualized SharePoint environment.

# Scope

This guide focuses on providing an understanding of the key considerations for virtualizing SharePoint Server 2013 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 Volumes.
- **Virtual machine configuration for SharePoint and SharePoint Virtual Topologies:** Highlights best practice considerations related to configuring virtual machines for SharePoint 2013. Provides information related to SharePoint 2013 resiliency in different scenarios, including a single VM and farms across multiple hosts.
- **System Center enhancements:** Provides an overview of how System Center 2012 SP1 supports deploying and managing SharePoint 2013 across the infrastructure (that is, on-premises and in the private cloud).

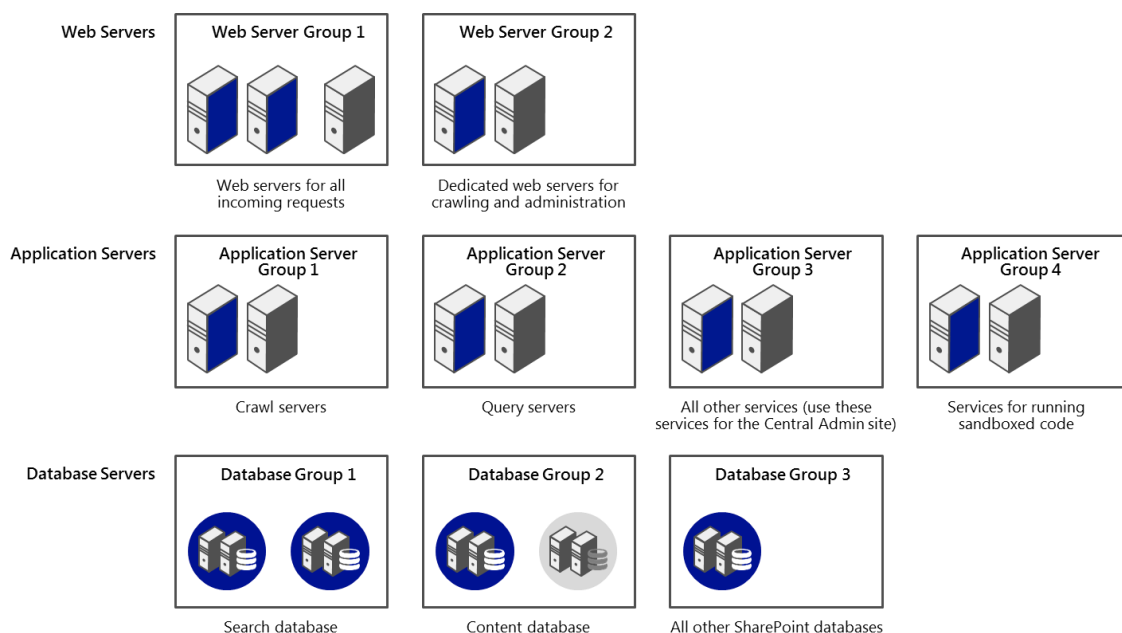
# Why Virtualize SharePoint?

Increasingly, organizations want to virtualize modern multi-tiered applications like SharePoint, to better meet their business and collaboration needs. According to a report from the Enterprise Strategy Group (ESG) Lab, among organizations already using virtualization in some way, approximately 53 percent are moving toward implementing virtualization technology for more complex and advanced systems.<sup>1</sup> For these organizations, it is necessary to consolidate and optimize computing resources for better flexibility, scalability, and manageability of mission-critical collaboration workloads like SharePoint 2013. This requirement is essential to better scale the key components of such demanding workloads—web servers, application servers, and database servers.<sup>2</sup>

In a traditional deployment of SharePoint, dedicated physical servers are usually used to deploy individual roles/components, including the front-end web server, application server, and database server (Figure 1). Organizations use separate physical servers for these roles to ensure high availability of services, better scalability, and improved performance. However, using separate physical servers for deploying separate roles has certain limitations, such as:

- **Underutilized resources:** CPU, memory, and storage are dedicated to a specific workload and remain idle while waiting for instructions, thereby consuming unnecessary power and space.
- **Higher costs:** Acquisition, maintenance, and management are more expensive.
- **Reduced efficiency:** A longer time is required to recover from outages. Plus, a higher Recovery Time Objective (RTO) may affect the service-level agreement (SLA).

Figure 1: Example SharePoint reference architecture (tiers and services)

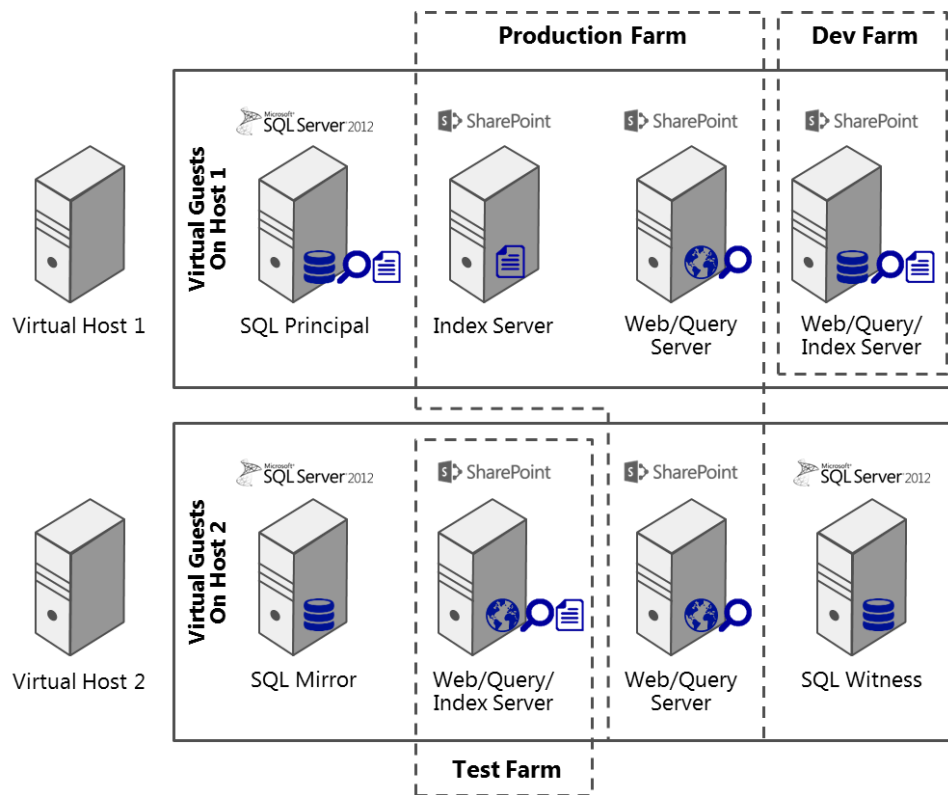


Deploying SharePoint on virtualization technology can help organizations gain multiple benefits, including reduced infrastructure and management costs, higher availability, and better management control. Returning to the report from the ESG Lab, 46 percent of organizations surveyed have already deployed

their corporate portals and collaboration solutions (such as SharePoint) on production virtual machines.<sup>3</sup> Virtualizing SharePoint provides the flexibility and rapid deployment capabilities that are required to meet complex and ever-changing business needs. It provides simplified administration and the ability to improve service levels by providing the ability to manage multiple virtual hosts, quickly provision SharePoint servers and farms, and migrate physical servers to virtual ones.

Virtualization enables organizations to “right-size” their hardware requirements and consolidate application workloads scattered on multiple underutilized physical servers onto fewer machines running virtualization technology (Figure 2).<sup>4</sup> With recent technological advances, complex applications can be consolidated more readily using virtualization. With the improvements in Windows Server 2012 Hyper-V, you also gain the same performance benefits when running complex applications on virtual servers as when running them on physical servers. Windows Server 2012 Hyper-V provides better consolidation of these workloads, which tend to saturate resources and contend for other system resources, including CPU, memory, and storage.

Figure 2: Example SharePoint virtualization farm architecture





# Why Microsoft Virtualization & 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 SP1, Windows Azure, and Microsoft Visual Studio 2012.

Windows Server 2012 Hyper-V is an optimal virtualization platform that can be used for deploying demanding and multi-tiered production applications, including key SharePoint 2013 components. With Hyper-V, Microsoft has become one of the leading vendors in virtualization technology.<sup>5</sup> 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.

By combining Windows Server 2012 with System Center 2012 SP1, organizations can comprehensively manage demanding applications (such as SharePoint 2013 workloads) as well as the infrastructure—including physical and virtual resources—in an integrated and unified manner.<sup>6</sup> The key benefits of this integrated virtualization and management platform by Microsoft include the following:<sup>7</sup>

- **Better scalability:** Higher capacity virtual machines supporting up to 64 vCPUs and 1TB memory per VM, and virtual machine density (up to 1,024 per host, and 8,000 per cluster).
- **Better performance:** Hyper-V support for Host & Guest Non-Uniform Memory Access (NUMA), Virtual Fibre Channel, Hardware Offloading, SR-IOV and more.
- **Better availability:** Faster and simultaneous live migrations, storage migrations and shared nothing live migrations, along with dynamic quorum for more resilient failover clusters.
- **Better manageability:** Same comprehensive management tools in System Center 2012 SP1, for SharePoint virtual machines in both private and public clouds.

# Fabric Configuration

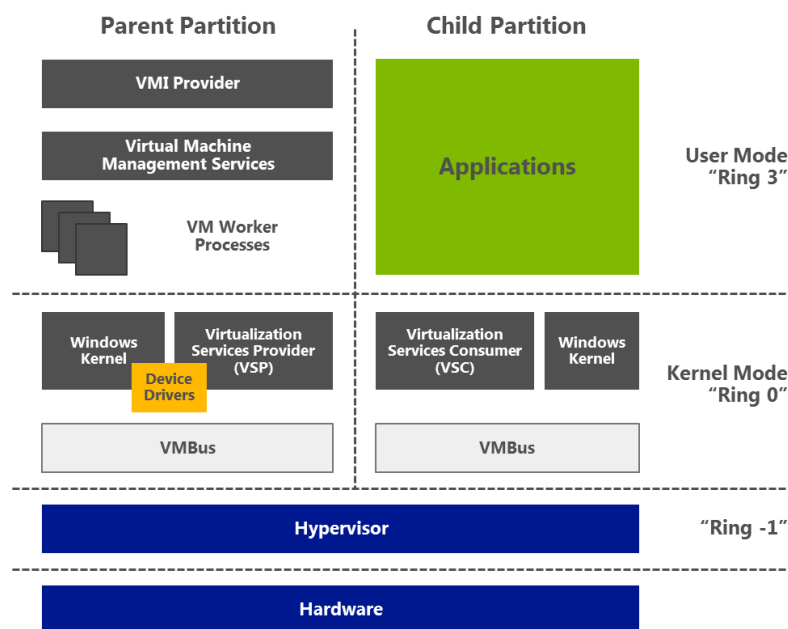
With Windows Server 2012 Hyper-V, organizations can take better advantage of the cost savings of virtualization and make the best use of server hardware investments by consolidating workloads as separate virtual machines. Windows Server 2012 provides a number of compelling capabilities to help organizations gain greater scalability and build reliable virtualized infrastructure for their mission-critical workloads like SharePoint 2013. 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 processors 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 3).

Figure 3: 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 minimum system requirements for Windows Server 2012 are as follows:<sup>8</sup>

- **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 recommend carefully considering the intended SharePoint 2013 workloads and their requirements when planning for hardware resources.

For an optimal experience and better performance and stability of Windows Server 2012, customers should utilize hardware that is Certified for Windows Server 2012. 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 multiple deployment options, including a Server Core Installation, Minimal Server Interface, and Server with a GUI.<sup>9</sup> 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. The *Minimal Server Interface* option in Windows Server 2012 does not include many aspects of Server Graphical Shell. With this option enabled, you can perform most of the GUI management tasks without requiring Internet Explorer or the full Server Graphical Shell.<sup>10</sup> Minimal Server Interface has more options/features than Server Core Installation, but it lacks the significant GUI components of the full installation. The *Server with a GUI* option is the Windows Server 2012 equivalent of the full installation option available in Windows Server 2008 R2.

#### Best Practices and Recommendations

Use the *Server Core Installation* option for setting up Hyper-V hosts in a SharePoint virtualization environment. This helps to reduce the servicing footprint and potential attack surface of the host. The Server Core installation option however, cannot be used to host the SharePoint 2013 components themselves – these should be installed on a full GUI installation of Windows 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.<sup>11, 12</sup>

Table 1: Resources available across versions of Windows Server

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

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).<sup>13</sup> 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 \(MAP\) Toolkit](#) provides data for Hyper-V virtualization planning; identifies server placements; and performs virtualization candidate assessments, including return on investment (ROI) analysis for server consolidation with Hyper-V.

## Compute Considerations

To fulfill scalability and high availability requirements, organizations need a virtualization technology that supports a large amount of processing and memory power. Windows Server 2012 Hyper-V greatly expands support for processors and memory to maintain massive-scale workloads. Other improvements enhance the availability and scalability of mission-critical workloads like SharePoint 2013 running on a virtualized infrastructure. Windows Server 2012 Hyper-V provides support for increased cluster sizes, a

greater number of active virtual machines per host, and more advanced performance features such as in-guest NUMA.

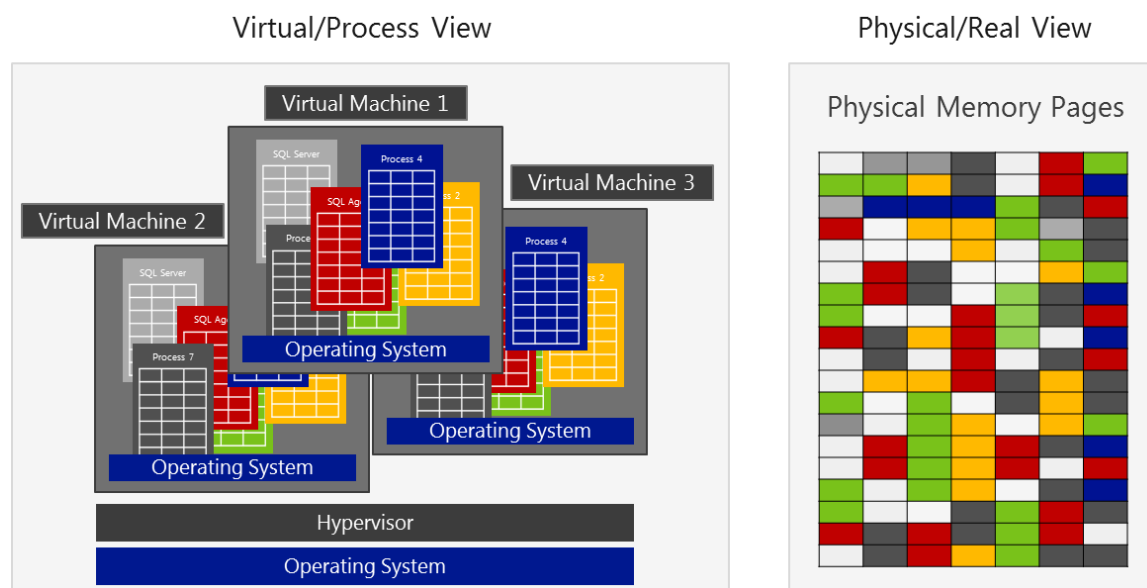
## Logical Processors on Hardware

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 the 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. Administrators 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. For using any virtual machine in a SharePoint 2013 farm, we recommend a ratio of 1:1.<sup>14</sup> Oversubscribing the CPU on the virtualization host can decrease performance, depending on how much the CPU is oversubscribed.

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.<sup>15</sup>

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 4).

Figure 4: 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, run Windows Server 2012 Hyper-V on SLAT-capable processors/hardware. This offers the additional benefits of improved performance, greater virtual machine density per host machine, and reduced overhead as compared to non-SLAT systems.

## 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.

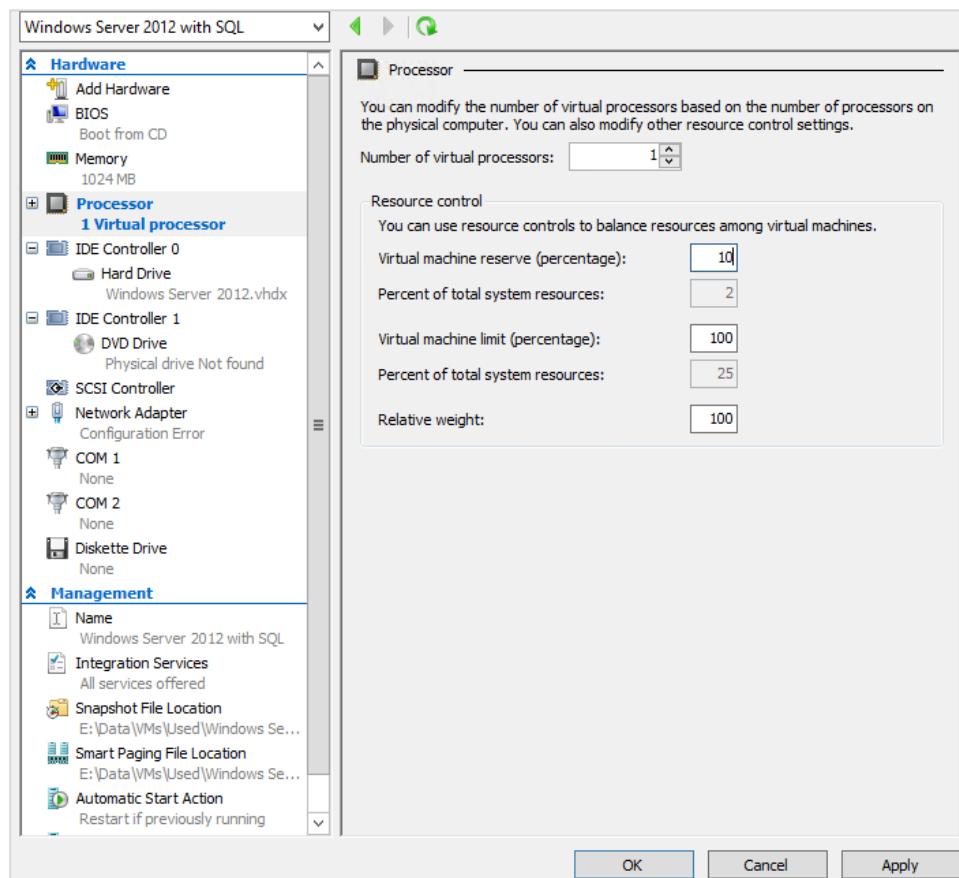
As previously discussed, Windows Server 2012 Hyper-V supports virtual machines with up to 64 virtual processors and 1 TB of memory. With increased support for 320 logical processors on a host machine/hardware, Hyper-V in Windows Server 2012 can now support up to 2,048 virtual processors per host. Windows Server 2012 Hyper-V supports massive scalability for even the most demanding workloads. Whether you are managing a vast enterprise content management deployment or addressing the need for greater collaboration due to acquisition, Windows Server 2012 provides better flexibility to scale up than previous versions of Windows Server. Plus, organizations do not have to worry about purchasing additional physical hardware resources.

### Best Practices and Recommendations

For the highest level of performance, configure a VP:LP ratio of 1:1 for any virtual machine that is used in a SharePoint 2013 farm. Remember that oversubscribing the CPU on the physical host used for virtualization can reduce performance.

Windows Server 2012 Hyper-V also provides the Weights and Reserves feature (Figure 5). 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.<sup>16</sup> 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 5: 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.<sup>17</sup>

NUMA is a memory design architecture that delivers 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 6). Each node has its own processors and memory, and is connected to the larger system through a cache-coherent interconnect bus.<sup>18</sup>

Figure 6: NUMA node (processor and memory grouped together)




Multiple NUMA nodes can exist in a host system (Figure 7). 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.<sup>19</sup>

Figure 7: 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.<sup>20</sup>

### 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.<sup>21</sup>

### 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. In the case of SharePoint environments, Microsoft does not support the use of Dynamic Memory, or technologies similar to Dynamic Memory found on alternative hypervisor platforms. This is because certain features of SharePoint can suffer from performance degradation when Dynamic Memory is enabled. For example, the cache size for the Search and Distributed Cache features are not resized when the memory allocated to the virtual machine is dynamically adjusted.

In most production SharePoint Server deployments, we recommend that you have at least 8 GB of RAM on each web server. Capacity should be increased to 16 GB on servers that have greater traffic or deployments with multiple application pools set up for isolation.

# Storage Considerations

Storage configuration is one of the critical components to any successful collaboration deployment. With an increased need for larger quotas, storage costs remain a concern in many organizations that have limited IT budgets to address storage needs. These organizations want high-capacity and low-cost commodity storage options, especially for larger content databases used only for document storage.

To optimize thinly provisioned logical unit numbers (LUNs), storage in a virtualized environment should be configured with properly sized I/O and exceptional storage of just-in-time allocations. This is achieved by maximizing the use of existing storage and reclaiming any unused space while taking into account thresholds and physical resource constraints. Optimizing LUNs provides sufficient I/O throughput as well as storage capacity to meet the current and future needs of most organizations.

## Storage Options for Hyper-V Virtual Machines

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.

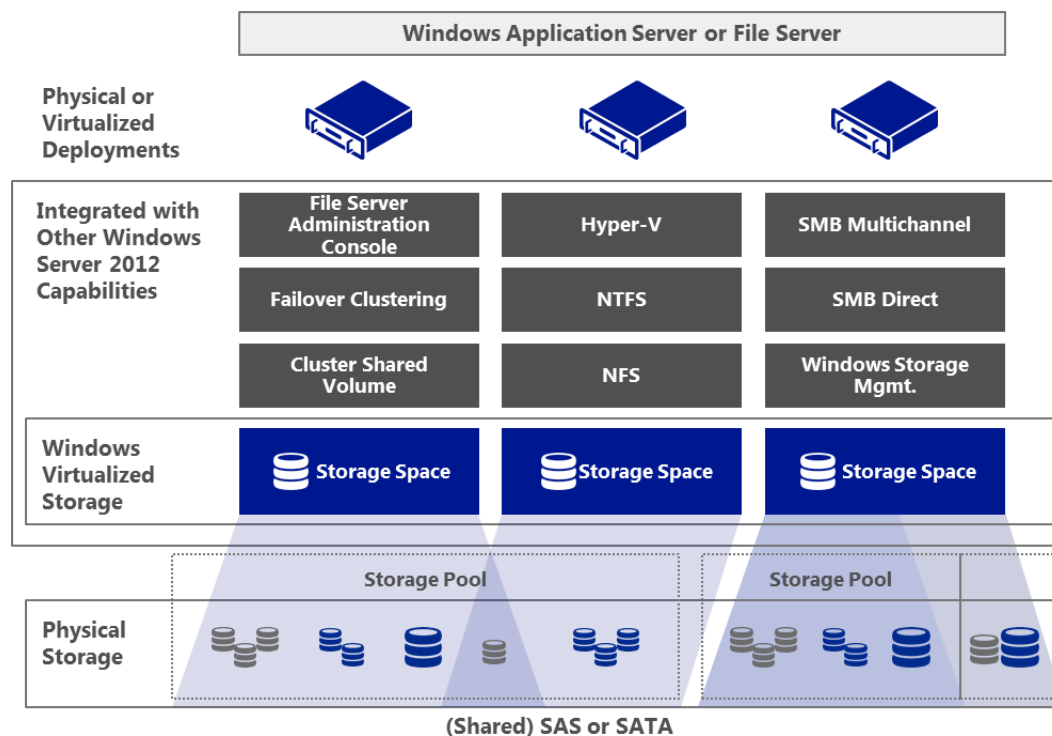
The Storage Spaces technology is fully integrated with failover clustering to enable continuously available service deployments. One or more storage pools can be clustered across multiple nodes within a single cluster. Storage Spaces supports thin provisioning to allow organizations to easily share storage 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 8). 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 8: 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.<sup>22</sup>

**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 9, there are four disks, and 1 MB of data needs to be written to disk. In this case, there are two options for writing data to disk: 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 9: 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 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 10, 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 10: 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 11, 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 11: Parity storage space across four disks

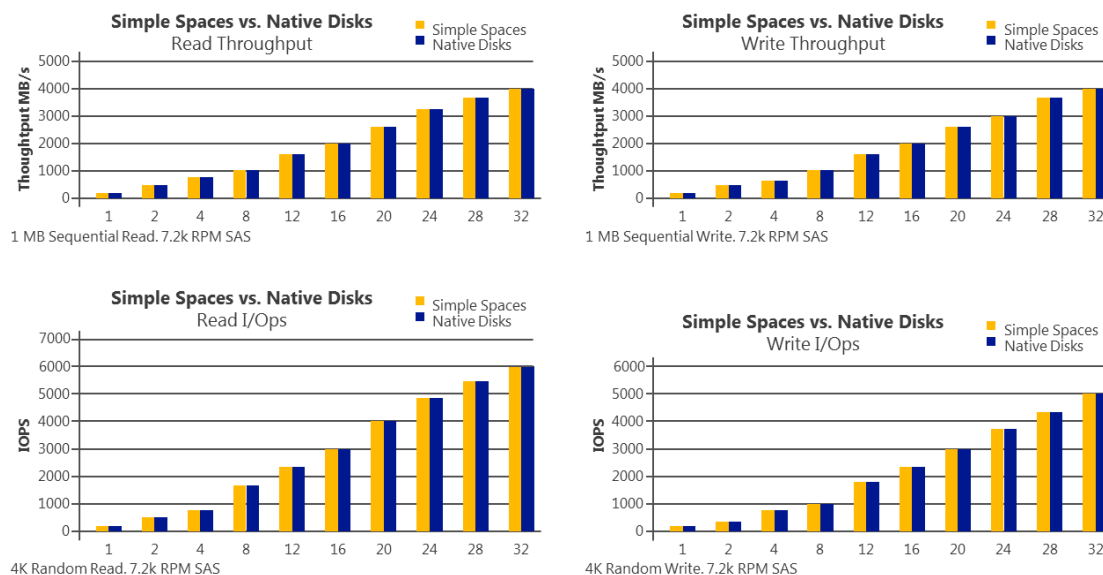


Parity storage spaces are used for the following:

- 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.

The graphs in Figure 12 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.<sup>23</sup>

Figure 12: Performance scaling of a simple storage space



## 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, and the Windows Server File Server also allows you to create a failover cluster without expensive storage area networks (SANs).

### Hyper-V over SMB

By enabling Hyper-V to use SMB file shares, you can greatly enhance performance 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 requires:

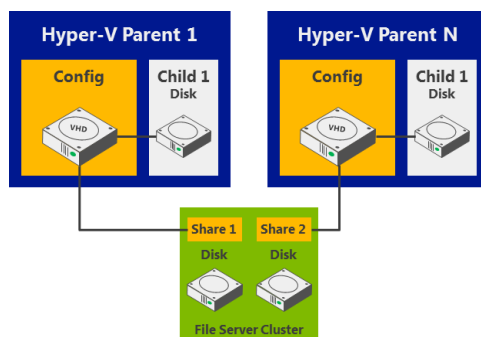
- One or more computers running Windows Server 2012 with the Hyper-V and File and Storage Services roles installed.
- A common Active Directory infrastructure. (The servers running Active Directory Domain Services do not have to run Windows Server 2012.)

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

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.<sup>24</sup>

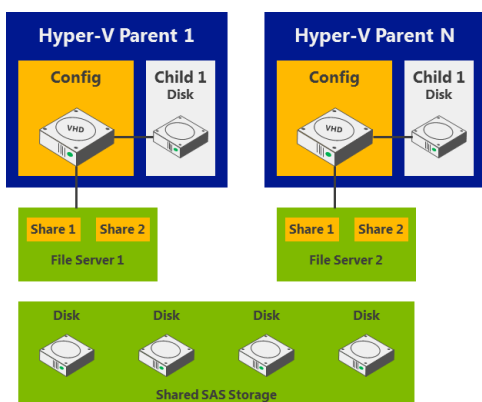
**Single-Node File Server:** In a Single-Node File Server, Hyper-V shares are used for VHD storage (Figure 13). File servers use standalone and local storage. This configuration provides flexibility for shared storage, as well as low costs for acquisition and operation. It does not provide continuous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

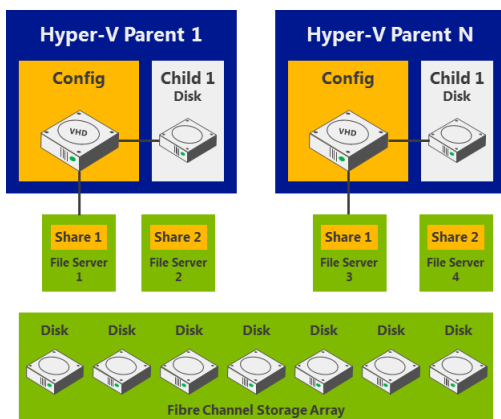


Table 2 compares the cost and availability/scalability of the three configurations for Hyper-V over SMB.

Table 2: Comparison of Hyper-V over SMB configurations

	Single-Node File Server	Dual-Node File Server	Multi-Node File Server
<b>Cost</b>	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)
<b>Availability/Scalability</b>	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 adapters to take advantage of SMB Multichannel as this provides increased performance and resiliency.

## 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 virtualized workloads on Hyper-V, which have virtual disks stored on SMB shares, SMB Direct provides a considerable boost to performance. 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).



## Best Practices and Recommendations


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

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



As you can see, as the number of Hyper-V virtual machines increased, the number of transactions per second increased whilst recording manageably low average transaction response times, even though the VMs and respective databases were residing on remote file servers, accessed using SMB 3.0. The full report, and further details are [available here](#).

## 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 Windows Server 2012 to act as an iSCSI array for other systems or workloads—with Windows Server 2012 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.

Windows Server 2012 also contains iSCSI initiator that allows Windows Server to quickly and easily connect to existing iSCSI storage and store key files, folders, and virtual hard disks etc. In particular, the iSCSI Initiator is the component in the overall iSCSI environment that creates the appearance that the iSCSI SANs are locally attached disks.

## 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.

Windows Server 2012 fully supports Fibre Channel connectivity for storage of virtual machine files. In addition, Windows Server 2012 Hyper-V provides a new capability for the virtual machines themselves, known as Hyper-V Virtual Fibre Channel (VFC). This allows connecting to Fibre Channel storage directly from within virtual machines opening up new scenarios around guest clustering, and a more direct path to the underlying FC fabric from within the virtual infrastructure.

## 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.

### Note

Some of the vendors supporting FCoE hardware include NetApp, Brocade, Cisco, Intel, QLogic, EMC, and Emulex.

FCoE requires switches that have data center bridging (DCB), which provides the extensions to traditional Ethernet that make it suitable for transporting storage traffic in a lossless way. DCB capability is available in some 10 GbE switches. Adapters that work with FCoE are known as converged network adapters (CNAs). Traditional Ethernet and Fibre Channel host bus adapter (HBA) vendors provide CNAs and support Ethernet and Fibre Channel simultaneously over the same wire. These CNAs run at 10 Gbps for both Ethernet and Fibre Channel.

## Note

Standard 10/100, 1 Gb or 10 Gigabit Ethernet (10 GbE) do not support FCoE. FCoE runs on versions of Ethernet that have been improved to provide low latency, quality of service, guaranteed delivery and other functionality traditionally associated with channel interfaces.

Fibre Channel, OM3 and OM4 cabling are suitable for FCoE and 10 GbE,


## Multipath I/O

Microsoft Multipath I/O (MPIO) is an in-box 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 up to 32 paths between storage devices and the Windows host operating system, and provides fault-tolerant connectivity to storage. Further, as 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.

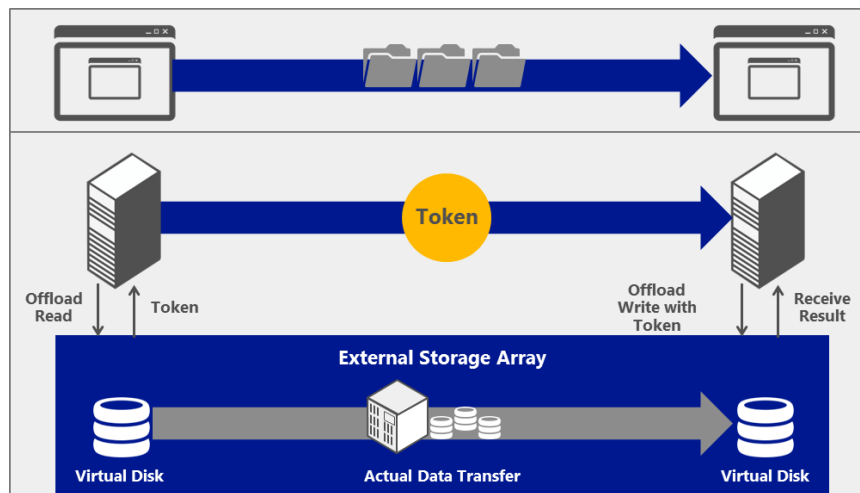
## Offloaded Data Transfer

For customers who have invested in storage technologies, such as iSCSI or FC SANs, Offloaded Data Transfer (ODX) in Windows Server 2012 enables you to accomplish more with existing external storage arrays by letting you quickly move large files and virtual machines directly between storage arrays, which reduces host CPU and network resource consumption.

ODX enables rapid provisioning and migration of virtual machines and provides significantly faster transfers of large files, such as database or video files. By offloading the file transfer to the storage array, ODX minimizes latencies, maximizes the use of array throughput, and reduces host resource usage, such as CPU and network consumption. File transfers are automatically and transparently offloaded when you move or copy files, regardless of whether you perform drag-and-drop operations in Windows Explorer or use command-line file copy commands. No administrator setup or intervention is needed.

To eliminate the inefficient and unnecessary steps required by traditional host-based file transfers, 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. The token serves as a point-in-time representation of the data. For example, when you copy a file or migrate a virtual machine between storage locations, Windows Server 2012 copies the token representing the virtual machine file. This removes the need to copy the underlying data between servers.

Figure 18: Offloaded Data Transfer in Windows Server 2012




ODX in Windows Server 2012 enables rapid provisioning and migration of virtual machines, lets you quickly move large files and virtual machines directly between storage arrays, and reduces host CPU and network resource consumption.

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.<sup>25</sup>

If you are using iSCSI, each clustered server should have a minimum of two 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 should 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 business applications connect to databases. Typically, a business 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 SharePoint 2013.

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

Network Interface Card (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 in-box 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 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.<sup>26</sup> With its new features and enhanced capabilities, the Hyper-V Extensible Switch supports tenant isolation, traffic shaping, protection against malicious 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 a number of top partners that offer networking extensions for Hyper-V environments, as well as their key extension products.

Table 3: Example Options for networking extensions

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

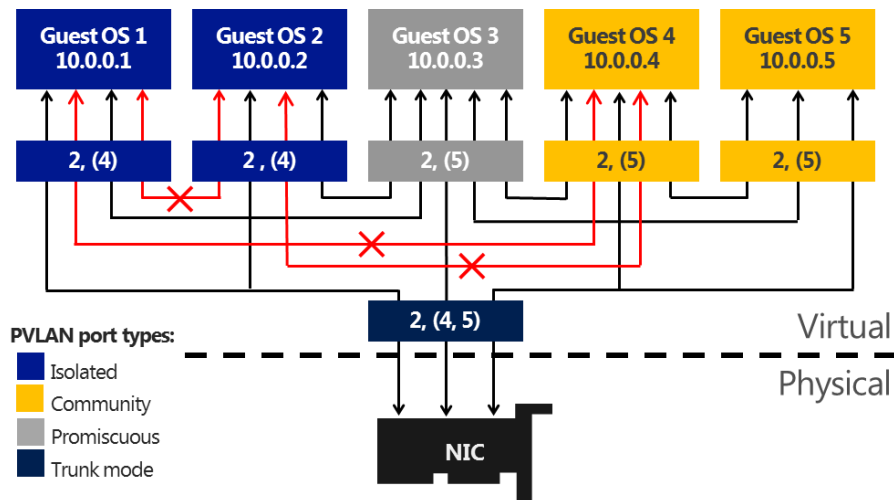
## Virtual LANs

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. 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. In 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

### Best Practices and Recommendations

VLANs and PVLANs can be a useful mechanism to isolate different SharePoint infrastructures, for instance, a Service Provider hosting multiple unrelated SharePoint farms, or a single enterprise, running separated farms for testing and development. For customers with VLAN constraints, PVLANs enable extra levels of isolation granularity, within the same VLAN. PVLANs can be configured through PowerShell.

## Hardware Offloads - Dynamic Virtual Machine Queue

Virtual Machine Queue (VMQ) allows the host's network adapter to pass DMA packets directly into individual virtual machine memory stacks. Each virtual machine device buffer is assigned a VMQ, which avoids needless packet copies and route lookups in the virtual switch. Essentially, VMQ allows the host's single network adapter to appear as multiple network adapters to the virtual machines, allowing each virtual machine its own dedicated network adapter. The result is less data in the host's buffers and an overall performance improvement to I/O operations.

The VMQ is a hardware virtualization technology for the efficient transfer of network traffic to a virtualized host operating system. A VMQ-capable network adaptor classifies incoming frames to be routed to a receive queue based on filters which associate the queue with a virtual machine's virtual network adaptor. These hardware queues may be affinity-tied to different CPUs thus allowing for receive scaling on a per-virtual machine network adaptor basis. Windows Server 2008 R2 allowed administrators to statically configure the number of processors available to process interrupts for VMQ. Without VMQ, CPU 0 would run hot with increased network traffic. With VMQ, the interrupts were spread across more processors. However, network load may vary over time. A fixed number of processors may not be suitable in all traffic regimes.

Windows Server 2012 on the other hand, dynamically distributes incoming network traffic processing to host processors, based on processor use and network load. In times of heavy network load, Dynamic VMQ (D-VMQ) automatically uses more processors. In times of light network load, Dynamic VMQ relinquishes those same processors. This ensures that the overall performance is increased, and ensures that VMs with D-VMQ enabled,

### Best Practices and Recommendations

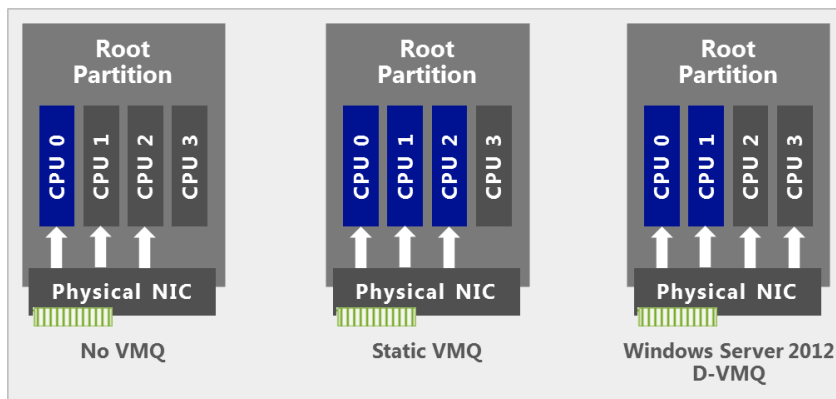
Some Intel multicore processors may use Intel Hyper-Threading Technology. When Hyper-Threading Technology is enabled, the actual number of cores that are used by D-VMQ should be half the total number of logical processors that are available in the system. This is because D-VMQ spreads the processing 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.)

VMQ provides improved networking performance to the **management operating system as a whole rather than to a specific virtual machine**. For the best results, treat queues as a scarce, carefully managed resource. Because queues are allocated to virtual machines on a first-come, first-served basis, making all virtual machines eligible for a queue may result in some queues being given to virtual machines with light traffic instead of those with heavier traffic. **Enable VMQ only for those virtual machines with the heaviest inbound traffic**. Because VMQ primarily improves receive-side performance, providing queues for virtual machines that receive the most packets provides the most benefit to overall management operating system performance

As shown in Figure 23, without the VMQ technology and RSS, the majority of network processing burdens CPU0 and ultimately limits the scale of the solution. With D-VMQ, processor cores are dynamically assigned to distribute the workload.

Figure 23: Dynamically distributed workload with D-VMQ for Hyper-V



# Host Resiliency & VM Agility

As organizations grow, effective collaboration becomes increasingly critical. This means that business data and sites must be highly available. Windows Server 2012 Hyper-V provides enhanced capabilities to ensure that SharePoint 2013 workloads are highly available and scalable, as well as resilient and agile.


## 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

All hardware should be certified for Windows Server 2012 and failover cluster solution should pass 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.

The aim of Failover Clustering is to provide a resilient solution for workloads that are running on top of the cluster. These could be clustered roles such as a File Server, a DHCP Server, or in our case, a Virtual Machine. Windows Server 2012 Failover Clustering provides resiliency for a number of other roles, or services, of which more information can be [found here](#).


With virtual machines running on top of a Hyper-V Cluster, should a Hyper-V cluster node fail, the virtual machines will also experience downtime, however the remaining cluster nodes are immediately working to bring those virtual machines up again on an alternative node within the cluster, ensuring that downtime is minimal, and administrator intervention is not required. This ensures that the workload, in this case, SharePoint 2013, has an extra level of resiliency when compared with a physical implementation with which the administrator can only rely on application-level resiliency, as oppose to dual layers of resiliency provided by a clustered Hyper-V implementation. This will be explored in more detail when we look at topologies later in the document.

## 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 LUNs in a failover cluster. CSV allows simultaneous read-write permissions to the same NTFS-provisioned LUN disk (Figure 26) from multiple cluster nodes. 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 25: 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. The Windows Server Catalog can be a useful source of guidance for this information.
- Isolate the LUN used for one set of cluster servers from all other servers through LUN masking or zoning.
- Use multipath I/O software for the highest level of redundancy and availability for when connecting your hosts to iSCSI or FC storage.

## Cluster Networking

Before you begin to construct your Failover Cluster, which will form the resilient backbone for your key virtualized workloads, it's important to ensure that the networking is optimally configured. Clusters, as a minimum, would require 2 x 1GbE network adapters, however for a traditional production Hyper-V Failover Cluster, it is recommended that a greater number of adapters is used to provide increased performance, isolation, and resiliency.

As a base level, for a customer utilizing gigabit network adaptors, it is recommended that a customer have **8 network adaptors**. This would provide the following:


- 2 teamed NICs that would be utilized for **Hyper-V host management** and **cluster heartbeat**
- 2 teamed NICs that would be utilized by a **Hyper-V Extensible Switch**, to allow **VM communication**
- 2 teamed NICs that would be utilized by **Cluster Shared Volumes** traffic and communication

- 2 teamed NICs that would be utilized for **Live Migration** of virtual machines

In addition, if a customer is utilizing **iSCSI storage**, there should be an **additional 2 NICs** added for connectivity from host to storage, and MPIO should be utilized instead of NIC Teaming. Alternatively, if the customer is utilizing **SMB storage** for the cluster, the hosts should have **2 NICs**, but SMB Multichannel (or SMB Direct, if RDMA is present) will be used to provide enhanced throughput and resiliency. So, for a customer with a need for a cluster that uses Cluster Shared Volumes, Live Migration, and a requirement for resiliency at the NIC level, this would present a requirement for 10 NICs as a minimum. These different networks can be combined onto fewer NICs, but isolation is recommended for optimal performance.

An alternative to this, would be the use of fewer, higher bandwidth NICs, such as 2 x 10GbE NICs, which would be combined in a host-level team, for an aggregated 20Gb bandwidth. The question is, how do I isolate all of the different types of traffic, such as CSV, Live Migration etc., on what essentially would be a single NIC team presented at the host level? To solve this problem, I could utilize a converged approach.

Figure 26: High level overview of cluster member converged networking configuration



In the above high-level example, the Hyper-V cluster node has 2 x 10GbE NICs, which are configured in a team, and for the isolated networks that are required, we can create Virtual NICs (vNICs) for the host OS. Each cluster node member uses a virtual network adapter (for Live Migration, for CSV, for Management etc.) to connect to the single Hyper-V Extensible Switch, which connects it to the physical network. Each tenant virtual machine is also connected to the same Hyper-V Extensible Switch using a regular virtual network adapter. Windows Server 2012 Hyper-V virtual switch QoS is used to assure that each traffic type (such as live migration, cluster, management and tenant) has a predictable amount of bandwidth available. Traffic isolation is enabled by 802.1q VLAN tagging so that host traffic is not visible to the tenants and Hyper-V virtual switch port ACLs can also be used for more granular access control at the network level.

This converged approach can significantly reduce the number of physical NICs required in each host, and subsequently, the number of overall switch ports, yet at the same time, still provides resiliency, and high levels of bandwidth for key VMs, and workloads. More information about the converged infrastructure options can be [found on TechNet](#).

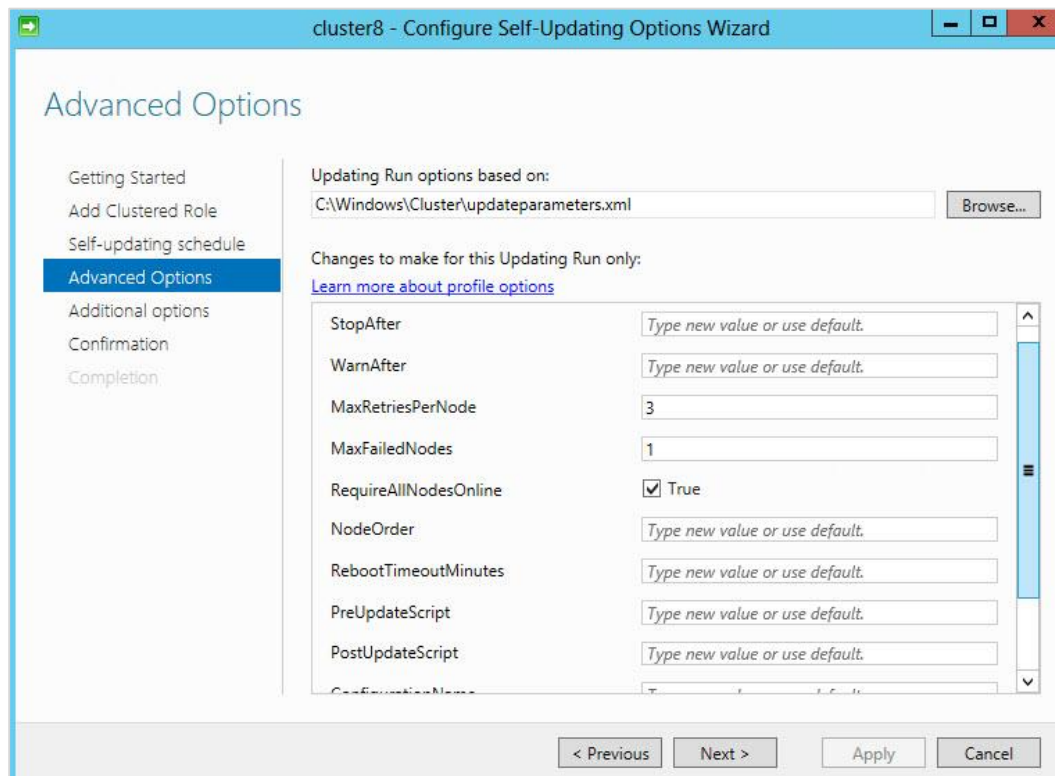
## Cluster-Aware Updating

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:<sup>27</sup>

- 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 27).

Figure 27: 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).<sup>28</sup>

### 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 virtual machines with live migration. This means that virtualized SharePoint servers will experience no downtime, even though underlying hosts will be patched, and potentially taken offline for maintenance.

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.

## 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, upon failover.<sup>29</sup>

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 the following:<sup>30</sup>

- **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

In case of failover, SQL Server virtual machines for SharePoint 2013 should typically be started first because they provide the back end for applications and services. The application server should be started next, then the web front ends. This will be discussed further when we look at some example topologies.

## 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, imagine two machines are partnered: One virtual machine has front-end applications and the other has a back-end database. It can be specified that two particular virtual machines cannot coexist on the same node in a failover scenario. This is the Virtual Machine Anti-Affinity 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 Anti-Affinity rules to keep related SharePoint virtual machines apart on hosts within a Hyper-V Cluster. This ensures that if any host is lost, whole SharePoint farms are not taken down as a result.

## Live Migration

We've spent a significant amount of time discussing the failover cluster, which provides the solid, resilient foundation for ensuring workloads, such as virtual machines, are as continuously available as possible. With a failover cluster in place however, certain other key capabilities, that can provide solutions for planned maintenance, are unlocked. Live Migration is one of those capabilities.


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, all without downtime. During the migration process of any workload, no additional configuration changes are required on the guest operating system.

- Windows Server 2012 with Hyper-V efficiently utilizes the available network bandwidth on your designated Live Migration networks, to reduce the time taken for the Live Migration.
- Hyper-V now offers functionality to migrate multiple virtualized SharePoint Servers simultaneously.

For migrating virtual machines, Hyper-V sets up a TCP connection between source and destination hosts to transfer virtual machine configurations. The memory assigned to the migrating virtual machine is transferred over the network to the destination 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 connectivity to physical storage attached through a virtual FC adapter. After completing all stages of migration, Hyper-V brings up the new destination virtual machine, with no downtime to the workload itself (Figure 28).

Figure 28: 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 several scenarios, including the need to redistribute virtual machine storage when disk space runs low or the need to perform maintenance on underlying storage without disrupting the virtual machine. Storage Live Migration for .vhd and .vhdx disks can be performed on standalone Hyper-V servers, and clustered Hyper-V Servers. Storage Live Migrations cannot be used with 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.

# 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

Consolidating an application's workloads can help to improve the efficiency and agility of that application, as well as provide better control and flexibility of computing resources in terms of their placement, sizing, and overall utilization. It is important to properly plan the consolidation of SharePoint workloads because every element of a virtual SharePoint Server farm needs to comply with [SharePoint licensing requirements](#). In addition, you must account for other software products that are prerequisites for installing SharePoint Server, including Windows Server and SQL Server.

The Microsoft Assessment and Planning Toolkit can help organizations to plan their SharePoint 2013 virtualization, speed migration to the virtual environment (Hyper-V), and start realizing the benefits of virtualization. The MAP Toolkit provides a complete software usage tracker. This tracker collects and reports client and server usage information of SharePoint deployments. The MAP Toolkit also provides multiple reports and proposals, including a server consolidation report, server consolidation proposal, workload discovery report (SharePoint, SQL Server, and so on), and cost savings and ROI assessment. The information provided in these reports and proposals can be used to consolidate SharePoint workloads and better utilize hardware resources.

## SharePoint Virtual Machine CPU Considerations

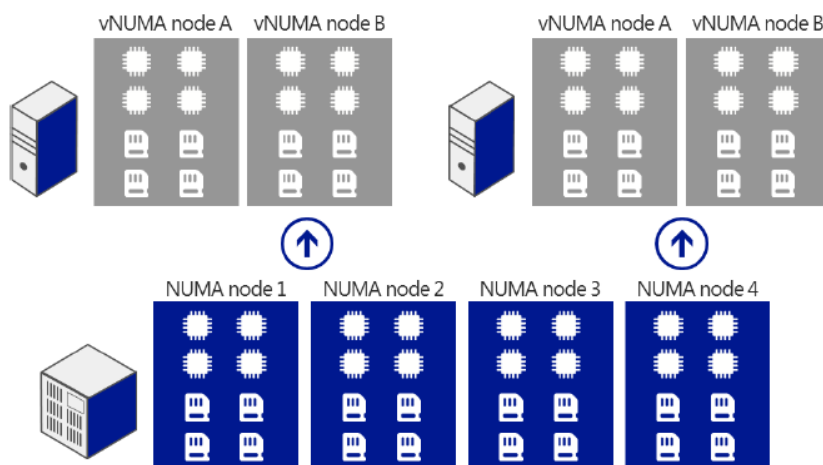
This subsection discusses Non-Uniform Memory Access from the virtual machine perspective. Virtual NUMA is the foundation for several key CPU considerations for SharePoint Server.

### 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 SharePoint Server support NUMA and use the computer's NUMA topology to increase performance by considering NUMA when scheduling threads or allocating memory. SharePoint 2013 is a multi-tiered application with back-end databases that perform many I/O-intensive transactions. Therefore, by reflecting the underlying NUMA topology within virtual machines, you can help to maximize the performance gains of running SharePoint 2013 virtualized farm environments on NUMA architecture.

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.<sup>31</sup>

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 SharePoint workloads.<sup>32, 33</sup> SQL 2012 and Internet Information Services (IIS) 8.0 on Windows Server 2012 are both NUMA-aware and provides the optimal configuration for the IT administrators.<sup>34</sup>

### 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 available performance counters.

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.

Do not oversubscribe the CPU on the virtualization host computer as it can decrease performance. For any virtual machine that you use in a SharePoint 2013 farm, use a virtual processor:logical processor ratio of 1:1 for optimum performance.

Identify and categorize virtual machines based on the intensity of the loads they bear (high intensity and low intensity). 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.

Install the latest virtual machine Integration Services in each supported guest virtual machine. Virtual machine Integration Services can help 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.

## SharePoint Virtual Machine 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<sup>35</sup>. One of these enhancements is known as Dynamic Memory, which allows Hyper-V to intelligently give and take memory from a virtual machine, whilst it's running, depending on the demand within that virtual machine at a given time.

### Best Practices and Recommendations

**Microsoft does not support Dynamic Memory for virtual machines that run any of the SharePoint 2013 components.** The reason is that this implementation of dynamic memory does not work with every SharePoint feature. For example, Distributed Cache and Search do not resize their caches when the allocated memory for a virtual machine is dynamically changed. This can cause performance degradation, especially when assigned memory is reduced.

## Dynamic Memory

Whilst not supported with the SharePoint 2013 components, Dynamic Memory can still provide a very valuable, and effective solution to optimizing memory usage for your other virtualized workloads. **VMs with Dynamic Memory enabled, can co-exist on hosts with other VMs that have Dynamic Memory disabled, without issue.** Earlier versions of Hyper-V only 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.<sup>36, 37</sup> 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 around how 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.

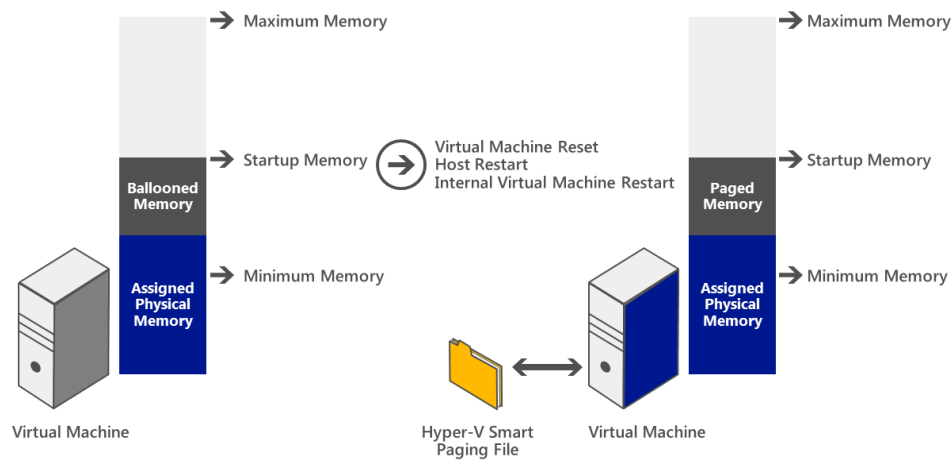
## Minimum Memory

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.<sup>38, 39</sup>

Figure 30: Hyper-V Smart Paging



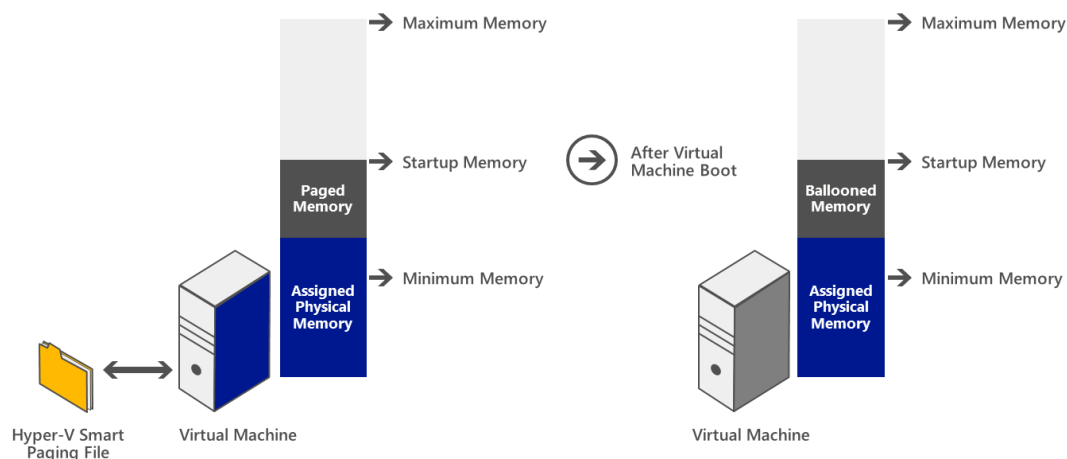
Hyper-V Smart Paging can lead to some performance degradation due to slower disk access speeds. Therefore, to ensure that the performance impact of Smart Paging is minimized, this feature is used only when all of the following are true:

- 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.

**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.<sup>40</sup>



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 SharePoint Server 2013 workloads.<sup>41, 42, 43</sup>

### Best Practices and Recommendations

Allocate a reasonable amount of memory to the virtual machines running SharePoint Server 2013 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 times or I/O usage for highly intensive SharePoint Server 2013 workloads. Also be sure to check the minimum memory requirement of the SharePoint Server components/roles that will be hosted on the Windows Server 2012 guest machine. Based on these two requirements, allocate the total minimum memory accordingly.

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

## SharePoint Virtual Machine Storage Considerations

With an optimal configuration of both CPU and memory, we need to ensure that the underlying disk subsystem is also optimally configured for the SharePoint workload. This subsection discusses two key storage considerations for SharePoint: virtual disks and guest storage.

### Virtual Disks

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

### VHDX File Format

Hyper-V in Windows Server 2012 introduces VHDX, a new version of the VHD, or Virtual Hard Disk 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 main new features of the VHDX format are:

- Support for virtual hard disk storage capacity of up to 64 TB.
- Protection against data corruption during power failures by logging updates to the VHDX metadata structures.
- Improved alignment of the virtual hard disk format to work well on large sector disks.
- Larger block sizes for dynamic and differencing disks, which allows these disks to attune to the needs of the workload.
- A 4-KB logical sector virtual disk that allows for increased performance when used by applications and workloads that are designed for 4-KB sectors.
- The ability to store custom metadata about the file that the user might want to record, such as operating system version or patches applied.
- Efficiency in representing data (also known as “trim”), which results in smaller file size and allows the underlying physical storage device to reclaim unused space. (Trim requires physical disks directly attached to a virtual machine or SCSI disks, and trim-compatible hardware.)

#### Best Practices and Recommendations

When creating virtual machines on Windows Server 2012 Hyper-V, the VHDX file format should be the default choice. Whilst not compatible with previous versions of Hyper-V, its capacity advantage, better alignment along with stronger protection against corruption, make it an ideal choice for key, mission critical workloads like SharePoint 2013.

## Dynamically Expanding or 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, with the VM offline, by using Hyper-V Manager or running a PowerShell script. Note that reducing the size is not supported. Fixed-size VHDX will deliver near-native-to-physical performance, and slightly higher performance than Dynamically Expanding VHDX files. Dynamic VHDX files, at creation time, only consume physical space based on their actual contents. For instance, an administrator could create a Dynamic VHDX with a maximum size of 127GB. Upon creation, the actual physical size of the VHDX file may only be a few MB, but as files are added to the VHDX file inside the Guest Operating System, the size of the VHDX file in the physical world, will grow accordingly. The Guest Operating System will always see the maximum size that the administrator chose upon creation.

#### Best Practices and Recommendations

Using differencing disks and dynamically expanding disks should be avoided in a virtualized SharePoint environment. The Thin Provisioned nature of the Dynamic VHDX file can mean that the underlying storage can become overcommitted, and as each Dynamic VHDX file grows in size, towards its configured maximum, the underlying storage could run out of space if not carefully

monitored. Instead, fixed-size VHDXs can be used to allocate a static amount of space on the underlying physical storage, up front, thereby ensuring that there will be enough storage space without unexpected surprises.<sup>44</sup>

To reduce disk contention, do not store system files on hard drives dedicated to storing virtual machines.

Do not use snapshots for the virtual machines in a SharePoint production environment. When you create a snapshot, Hyper-V creates a new secondary drive for the virtual machines. Write operations occur on the new drive, and read operations occur on both drives, resulting in reduced performance.

Be aware of underlying disk read/write contention between different virtual machines and their virtual hard 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 with a boot VHD/x file, because the virtual SCSI disks requires a driver to be present during boot-up. This driver is only present when booted into the OS. 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

The Virtual IDE controller has to be used for booting the virtual machine, however all other drives should be attached to the Virtual SCSI controller. This ensures the highest performance, but also the most flexibility. Each virtual machine has a single Virtual SCSI controller by default, yet 3 more can be added whilst the VM is offline.

## Guest Storage

In addition to presenting VHD or VHDX files to the respective SharePoint VMs, the administrator can choose to connect the Guest OS of the SharePoint VM directly to existing storage investments. Two methods that are provided in Windows Server 2012 Hyper-V, are using In-Guest iSCSI, and Virtual Fibre Channel.

### In-Guest iSCSI

Deploying a SharePoint 2013 virtual machine on iSCSI storage provides a cost-effective solution for enterprise-level virtual machine deployments. Instead of using virtual disks, such as the VHD or VHDX files discussed earlier, and placing them on the iSCSI LUNS presented to the host, the administrator can choose to bypass the host and connect the VMs directly to the iSCSI array itself. The iSCSI target, part of the storage array itself, provides storage to the SharePoint virtual machine directly over the VM's network

adaptors; the SharePoint VM uses the in-box iSCSI initiator inside the Windows Server Guest OS to consume the storage over a vNIC that has connectivity on the iSCSI storage network. The respective SharePoint servers can therefore store their information, logs, and other critical data directly on iSCSI disk volumes.

To enable this, administrators will need to create dedicated Hyper-V virtual switches and bind them to appropriate physical NICs in the hosts, to ensure that the VMs can communicate with the iSCSI storage on the appropriate network/VLAN. Once configured, the administrator must use the Guest OS IQN from the iSCSI initiator to present the appropriate LUNS directly to the VM, over the virtual networks. vNIC features such as Jumbo Frames, and some of the other offload capabilities can help increase performance and throughput over the network also. It's important to note that if you are intending to run the VM with In-Guest iSCSI, on top of a Hyper-V Cluster, all cluster nodes must have the same iSCSI virtual switches created on the hosts, to ensure when the VM migrates around the cluster, connectivity to the underlying storage isn't lost.

For resiliency, the administrator may want to use multiple vNICs to connect the VM to the iSCSI SAN. If this is the case, it's important to enable, and configure MPIO, as discussed earlier, to ensure optimal performance, and resiliency.

In-Guest iSCSI can also be used to create Guest Failover Clusters, which for a SharePoint environment, could be very useful when creating a SQL Server 2012 AlwaysOn Availability Group or Failover Cluster Instance, as a resilient backend for the SharePoint databases.

## Virtual Fibre Channel

In a similar way to the use of iSCSI that we discussed earlier, Virtual Fibre Channel for Hyper-V helps to connect to FC storage from within a virtual machine, bypassing the host's OS. It provides direct SAN access from the guest operating system by using standard World Wide Node Name (WWNN) and Worldwide Port Names (WWPN) 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 connected to the underlying, shared FC storage. As mentioned with iSCSI, this would be useful for deploying SQL Server 2012 AlwaysOn Availability Groups or Failover Cluster Instances, as a resilient backend for SharePoint 2013 databases.

For virtualizing SharePoint 2013, Virtual FC for Hyper-V allows you to use existing FC investments to drive the highest levels of storage performance access, while also retaining support for virtual machine live migration and MPIO.

## SharePoint Virtual Machine Networking Considerations

Networking, and network access is critical to the success of a SharePoint deployment. There are a number of capabilities, technologies and features within Windows Server 2012 Hyper-V that an administrator can take advantage of in order to drive the highest levels of networking performance for the virtualized SharePoint farm.

## Legacy versus Synthetic Virtual Network Adaptors

When creating a virtual machine, the administrator has 2 choices of virtual network adaptor, or vNIC; Legacy, or Synthetic. A legacy adapter emulates an Intel 21140-based PCI Fast Ethernet Adapter, which results in a lower data transfer than the network adapter. Legacy Network Adapters (also known as Emulated NIC drivers) should only be used for Pre-Boot Execution Environment (PXE) booting a VM or when installing non-Hyper-V aware Guest operating systems.

Synthetic adapters are the preferred option for most virtual machine configurations because they use a dedicated VMBus to communicate between the virtual NIC and the physical NIC. This results in reduced CPU cycles, as well as much lower hypervisor/guest transitions per operation. The driver for the synthetic adapter is included with the integration services that are installed with the Windows Server 2012 guest operating system<sup>45</sup>.

### Best Practices and Recommendations

From a SharePoint perspective, there should be no reason to use the Legacy vNIC. As a minimum, customers should be using the default Synthetic vNIC to drive higher levels of performance, and in addition, should the physical network card support them, take advantage of a number of the NIC offloads that can increase performance further.

## Single Root I/O Virtualization

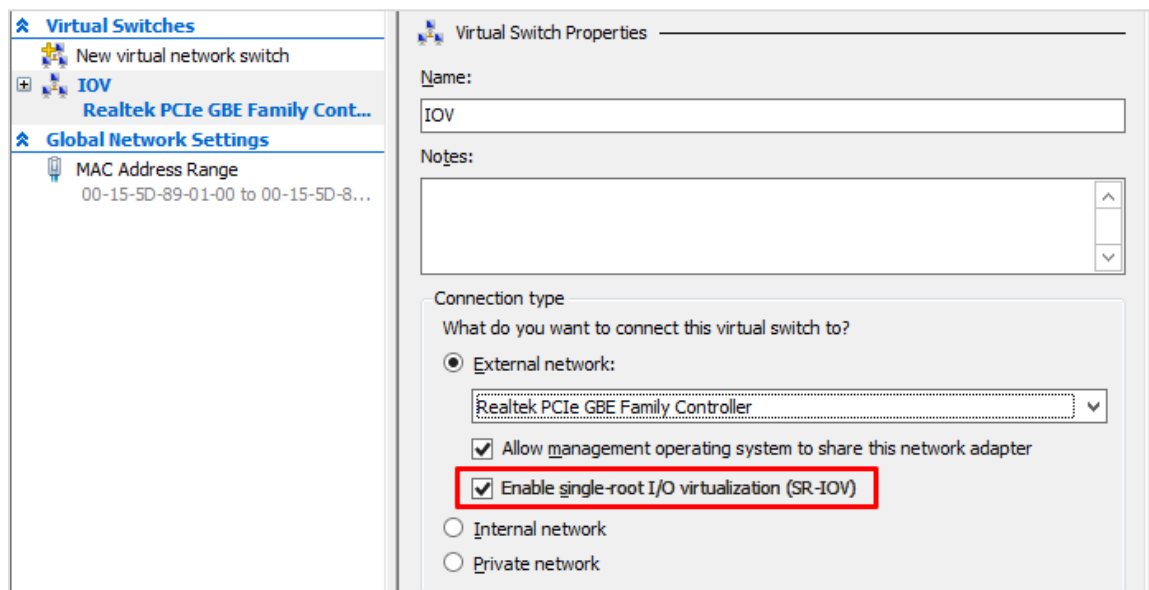
The Single Root I/O Virtualization standard was introduced by the PCI-SIG, the special interest group that owns and manages PCI specifications as open industry standards. SR-IOV helps to virtualize demanding workloads like SharePoint Server that require high network and I/O performance. It does so by 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 32).<sup>46</sup>

Figure 32: 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 33).

Figure 33: Enabling SR-IOV in the Virtual Machine Properties window

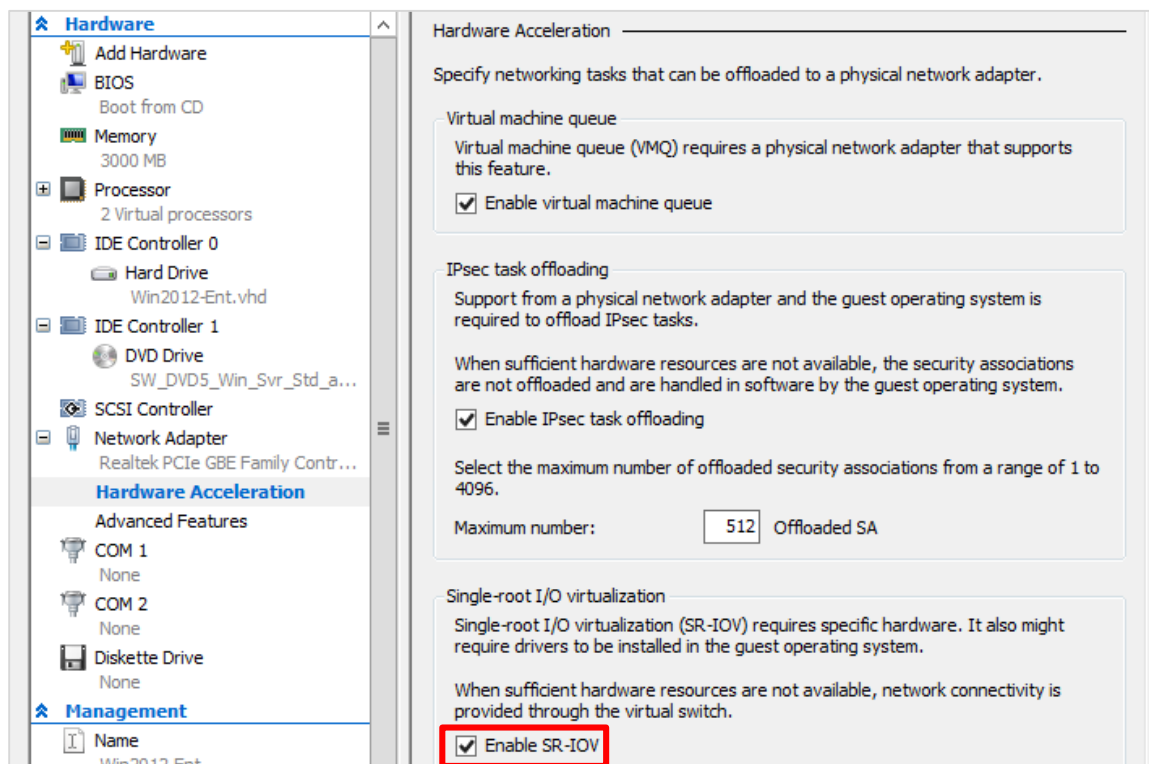


Figure 34 shows how SR-IOV attaches a physical NIC to a SharePoint Server virtual machine. This provides the SharePoint VM with a more direct path to the underlying physical network adaptor, increasing

performance and reducing latency, both of which are very important considerations for the SharePoint workload.

With the VM now attached to the physical NIC, through the Virtual Function, does this now mean that this particular virtual machine cannot be migrated, live, to another physical host? The answer is no, the VM is still free to be migrated, live, with no downtime, to another available node in the cluster, ensuring that Hyper-V in Windows Server 2012 is not only providing high levels of performance, but it's doing it without sacrificing agility.

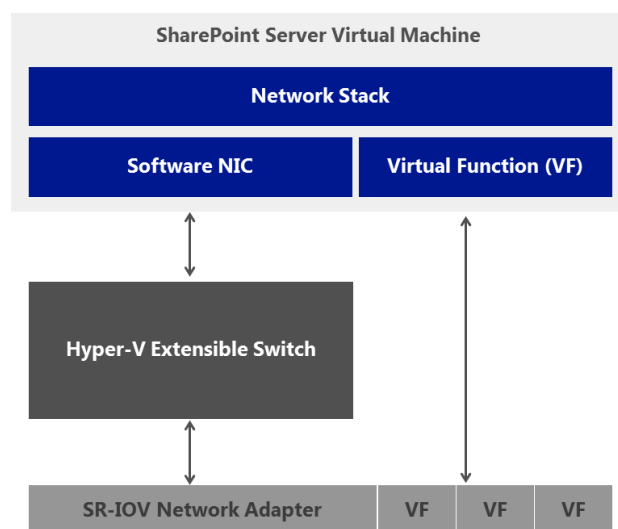
How does it do this? Well, SR-IOV addresses this concern by failing over the network traffic from the Virtual Function (VF) to the synthetic data path of the Hyper-V VM, automatically. The transition between the VF and synthetic data paths occurs with minimum loss of packets and prevents the loss of TCP connections. So whenever there is a state transition that would require hardware state to be saved, the VFs are removed from the VM beforehand, falling back to the synthetic path. Once the VF is removed, any operation necessary can be performed on the VM because it is a complete software-based container at that point.

The Hyper-V child partition is being live migrated to a different host at this stage. Once the operation has been completed, assuming hardware resources are available and other dependencies met, the VF is returned back to the VM. This solves the problem of saving the hardware state of the VM and helps in ensuring workloads receive the highest levels of performance.<sup>47</sup> This process looks like:<sup>48 49</sup>

- SR-IOV is enabled for the SharePoint virtual machine, and its VF is assigned.
- A team is automatically created inside the SharePoint virtual machine, and the traffic flows down the VF path, not the software stack.
- When live migration is initialized, the team is broken and the VF is removed. Traffic fails over to the synthetic path inside the VM.
- At this point, live migration of the SharePoint 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. Alternatively, if the VM has been migrated to a new host that doesn't have SR-IOV capable hardware, the network traffic will continue to operate along the synthetic software stack.

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

Figure 34: SR-IOV support in Hyper-V



If the SharePoint Server virtual machine is configured to use SR-IOV, but the guest operating system does not support it, SR-IOV VFs are not allocated to the virtual machine. We recommend that you disable SR-IOV on all virtual machines that run guest operating systems that do not support SR-IOV.<sup>50</sup>

### Best Practices and Recommendations

**SR-IOV can provide the highest levels of networking performance for virtualized SharePoint Servers.** Check with your hardware vendor for support because there may be a BIOS and firmware update required to enable SR-IOV.

## 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 SharePoint 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.<sup>51</sup>

For Hyper-V VMs specifically, QoS bandwidth management helps to set a throttling rate for a workload like SharePoint 2013. Both Minimum Bandwidth and Maximum Bandwidth enable organizations to enforce predictable network throughput the SharePoint workload. Apart from bandwidth management, organizations can prioritize and tag traffic so that QoS is enforced from end-to-end across a data center.

## Other Key SharePoint VM Considerations

In Windows Server 2012, Hyper-V Integration Services (IS) include six components that provide performance enhancements to a child partition (i.e., virtual machine or guest), and additional interoperability between child and parent partitions. Integrations Services are available in a child partition



only after they are installed in a supported guest operating systems. It is also possible to update Integration Services after the initial installation, and this is usually recommended when migrating a virtual machine from an older to a newer version of Hyper-V (e.g., Windows Server 2008 R2 to Windows Server 2012), or as new versions of the Integrations Services are released.

#### Best Practices and Recommendations

Running a Guest OS of Windows Server 2012 on a Windows Server 2012 Hyper-V host **will not require an update** of the integration services – they will automatically have the latest, most optimized versions.

Integration Services are installed as user mode components in the guest operating system, and are implemented in the following services:

- Hyper-V Heartbeat Service (vmicheartbeat)
- Hyper-V Guest Shutdown Service (vmicshutdown)
- Hyper-V Data Exchange Service (vmickvpexchange)
- Hyper-V Time Synchronization Service (vmictimesync)
- Hyper-V Remote Desktop Virtualization Service (vmicrdv)
- Hyper-V Volume Shadow-Copy Requestor Service (vmicvss)

Integration Services in a child partition communicate over a Virtual Machine Bus (VMBus) with components in the parent partition virtualization stack that are implemented as virtual devices (VDev). The VMBus supports high-speed, point-to-point channels for secure inter-partition communication between child and parent partitions. A dedicated VDev manages each of the parent partition Integration Services function, just as each dedicated service manages the different Integration Services functions in a child partition. Through this architecture, Integration Services components provide enhanced functionality and performance for mouse, keyboard, display, network, and storage devices installed in a virtual machine.

#### Best Practices and Recommendations

Disable the time synchronization for each SharePoint virtual machine. SharePoint 2013 implements timer jobs extensively and the latency during time synchronization will cause unpredictable results in the SharePoint environment. Instead, ensure that your virtual machine Guest OS's retrieve their time from an authoritative time source, such as a domain controller.

Further detailed information on the Integration Components, is [available on the TechNet Wiki](#).

For each virtual machine you can configure automatic stop and start behavior if a physical computer shuts down. The options for stop are as follows:

- **Save State** - The current state of the virtual machine is saved. When the virtual machine is started, Hyper-V attempts to restore the virtual machine to the state it was in.
- **Turn Off** - This is the equivalent of pulling the power plug on a server.

- **Shut Down the Guest OS** - This is the equivalent of shutting down a computer by using the Windows Shut down option.

#### Best Practices and Recommendations

For a SharePoint virtual machine, do not configure the virtual machine to save state. Virtual machines that start from saved state will be out of synchronization with the other servers in the farm. We recommend that you configure the virtual machine to use a shutdown because it minimizes the chances that the virtual machine can be corrupted. When a shutdown happens, all timer jobs that are running can finish, and there will be no synchronization issues when the virtual machine restarts.

The opposite of an automatic stop is an automatic start. Hyper-V provides the following startup options when the physical server restarts:

- **Do nothing** - You have to start the virtual machine manually regardless of its state when the physical server shut down.
- **Automatically start** if the virtual machine was running when the service stopped.
- **Always start** this virtual machine automatically - Hyper-V starts the virtual machine regardless of its state when the physical server shut down.

#### Best Practices and Recommendations

We recommend that you select either of the first two options. Both options are acceptable. However, the decision is ultimately up to the IT team that manages and maintains the virtual environment. In addition to the previous start options, you can configure a startup time delay for a virtual machine. We recommend that you do this to reduce resource contention on a virtualization host. However, if your start option is to do nothing, this is not an issue.

# SharePoint Virtual Topologies

To ensure the appropriate level of SharePoint availability in your organization, it is necessary to design the right solution for SharePoint 2013 on virtual infrastructure. High availability and disaster recovery are subsets of a broader Business Continuity Management (BCM) strategy and process that focus on the immediate design elements required to put a farm into production. The primary objective of a high availability and disaster recovery solution is to minimize the impact of downtime and reduce or eliminate data loss. To optimize availability and performance, you need a highly available topology and sound system requirements for a SharePoint Server farm.

However, before we design a virtualized SharePoint infrastructure, you must be familiar with the supported logical components of a SharePoint farm, such as web applications, service applications, and site collections. You must also be familiar with physical components, such as web servers and database servers so that you can determine the appropriate architecture for your SharePoint environment. For more information on these fundamentals, see the [Plan logical architectures for SharePoint 2013](#), the [Topologies for SharePoint Server 2013](#) model, the [Services](#) model, and the [Architecture design for SharePoint 2013 IT Pros](#) Resource Center on TechNet.

We recommend that you approach design of a virtual farm in the same way that you design a physical farm. Most issues and requirements to deploy on physical servers apply equally to virtual machines. Any decisions that you make, such as minimum processor or memory requirements, have a direct effect on the number of servers in the farm topology and corresponding virtualization host requirements. The host computers must have sufficient resources to support the virtual machines that you identify for the farm.

The following list shows some typical farm requirements that apply to physical and virtual farms.

- Identify the required roles (for example, queries) for a specific SharePoint solution.
- Decide how to distribute the SharePoint roles on farm servers.
- Determine how many servers are needed for the farm.
- Decide how each server needs to be configured

The starting point to determine farm server specifications are the minimum hardware and software requirements for SharePoint 2013. For more information about these requirements (such as specific updates that you must install), see [Hardware and software requirements for SharePoint 2013](#).

If your strategy is to implement a heterogeneous SharePoint virtualization environment, then you have to identify servers that are the best candidates for virtualization. From a technical and Microsoft support perspective, all servers that run SharePoint 2013 can be virtualized. The decision to virtualize a particular farm server should be based on specific requirements and considerations, such as the following:

- Corporate compliance policies (for example, legal and technical) might not support virtualization, or virtualization host server and storage location.
- Estimated performance and capacity requirements.
- Security requirements.

- Support and maintenance requirements

After you design the topology to support the farm, you have to understand farm virtualization and the supported virtual architectures. Architects who plan a virtual environment often apply the requirements for a SharePoint solution that was deployed in a physical environment to the virtual environment. This includes elements such as farm topology, capacity and performance requirements, and business continuity. You can use the number of database servers, front end web servers, and application servers — including specific application server roles — to determine the number of virtual machines and virtualization host server computers that a virtualized farm requires.

## SharePoint Products

If you have an existing SharePoint farm, you can use historical benchmark data and usage profiles as a starting point to determine virtual machine configurations for the different server roles in the farm. For example, concurrent connections, types of requests, and peaks in demand are useful types of data. After you have the virtual machine requirements, you then determine the number and capacity of the host servers that you require to support the virtual farm servers.

If you are upgrading an existing farm, remember that you have to have new benchmark data to validate your assumptions. This is a requirement because customers typically clean up and architect their SharePoint solutions again as part of the upgrade, which we recommend. Another reason to collect new data is that you might add functionality to the farm. SharePoint 2013 also provides new features and redesigned features, such as Search.

For new deployments, you have to base your starting requirements on the SharePoint 2013 planning, design, and sizing documents that are available in the Microsoft TechNet Library. After that you have to test, resize, and retest until you feel comfortable with your farm design.

## Identify Farm Server Roles to Virtualize

Identifying the server roles that you want to virtualize is related to the broader virtualization goals and objectives of your organization. Clearly, if the goal is implement an architecture where all the SharePoint farm servers are virtual, then the pros and cons of server selection is not an issue. In a homogeneous environment, the criteria for the design of the architecture are based on operational requirements such as performance and capacity and the host infrastructure that is needed to support the farm.

Although server virtualization is fully supported for all the server roles of SharePoint 2013, the server role alone is not the determining factor in deciding whether to virtualize some farm servers or all of the farm servers. Several factors will affect your virtualization strategy and architecture: IT constraints, host server capacity, and operations.

### IT Constraints

Many organizations do not support or allow IT departments to virtualize database servers. This policy is typical in organizations that have dedicated database teams that tightly manage and maintain SQL Server. In these controlled deployments, the database team has to create all databases. Virtualizing SQL Server is not an option.

### Host Server Capacity

The availability of host hardware than can adequately support the requirements of all the roles is something else to consider. For example, the CPU and memory requirements of a virtual database server are greater than a web server or application server. The issue of host hardware capabilities happens when you re-purposing existing hardware as part of a virtualization strategy.

## Operations

Managing and maintaining a virtual farm environment is complex and requires specific skills and tools. You have to manage the farm on the virtual level and the physical host level. If you deploy SharePoint 2013 in a partly virtualized environment, you have to deal with the virtual machines and their hosts. Additionally you have to maintain the physical computers that are used for specific roles. The result is that you must have three sets of skills, tools, and procedures to support the farm. **Full virtualization reduces the support requirements and simplifies support for the operations team.**

If there aren't any constraints about virtualizing specific server roles, and you don't know whether to virtualize some or all of the servers, the next step is to review all of the server roles and decide which ones to virtualize.

## Virtualizing Web Servers and Application Servers

The front end web server and application server roles are usually the first choices for virtualization.

The web server role is usually the first choice for virtualization because the workload demand on these servers is usually much lower than on other servers in a farm. As a result you can configure the virtual machine for this role to use fewer processors, less memory and fewer (and smaller) hard disks. The smaller resource footprint means that you can deploy more front end web servers on a single host than a highly specialized and resource-intensive system such as a database server, and in some cases, an application server. For organizations that are just starting to move to virtualization, virtual machines for the web server role are easier to plan for than the other roles, have the lowest virtualization host requirements, and are perceived as having the lowest risk in a production environment.

Application servers are also good initial candidates for virtualization. Depending on the degree of specialization, which is reflected by services they provide, they do not always have low resource requirements. A good example is an application server that hosts the search crawl component.

## Virtualizing SQL Server

IT professionals vigorously debate whether to virtualize SQL Server. Until recently the conventional wisdom was to avoid virtualizing the database server. As is the case for other farm servers, hypervisor imposes a performance cost. You have to define your performance goals and collect benchmark data before you decide whether to virtualize SQL Server. However, with Windows Server 2012 Hyper-V, and the increased scalability and performance capabilities, customers should feel confident in virtualizing tier-1, mission-critical, intensive workloads like SQL Server. Enterprise Strategy Group [recently performed](#) performance testing for SQL Server 2012 on Windows Server 2012 Hyper-V, and found that with Hyper-V's support for up to 64 vCPUs, ESG Lab took an existing SQL Server 2012 OLTP workload that was previously vCPU limited and increased the performance by six times, while the average transaction response times improved by five times, and in addition, saw a manageably-low Hyper-V overhead of 6.3% was recorded when comparing SQL Server 2012 OLTP workload performance of a physical server to a virtual machine configured with the same number of virtual CPU cores and the same amount of RAM.

Microsoft has also released [detailed guidance on virtualizing SQL Server 2012](#) on Hyper-V, to ensure that other workloads that rely on SQL Server, such as SharePoint 2013, can benefit from that configuration.

It's important that when virtualizing your key workload that you perform the appropriate performance testing within your environment, to ensure that the infrastructure behaves as expected. There are many different configurations, and topologies that you could deploy with SharePoint 2013, across any number of Hyper-V hosts, however we will discuss a small number, and some of the key considerations around each of the configurations. For further information, [see the TechNet library](#).

A SharePoint 2013 topology includes three-tier solutions and farm scenarios that may include a web-based front end, application servers, and database servers. This can be deployed on a single server for evaluation or development, or on many servers.

This section describes some of the key SharePoint topologies from the perspective of running on a resilient fabric (cluster) as defined in earlier sections:

**One Server Farm:** In this topology, all the roles (i.e. Web server role, Application server role and the Database server role) are installed on a single server. This type of topology is typically used for product evaluation, development and testing, or for environments that have limited numbers of users, say, less than 100, and don't require fault-tolerance.

Figure 35: One Server Farm in a VM




From a VM perspective, if this virtual machine, containing all 3 of these tiers, is running on a single Hyper-V host, not only would the VM not be resilient at the SharePoint level, but in addition, it would not be resilient at the host level either, meaning that if the VM failed, or the host failed, the entire SharePoint infrastructure would be down.

That said, even with this small scale deployment, by harnessing some of the capabilities we've discussed, such as taking advantage of **Virtual NUMA architecture**, higher performance **SR-IOV networking**, the increased capacity and performance of the **Fixed VHDX** file format, and the **enhanced scalability** of Hyper-V virtual machines in Windows Server 2012, an administrator could still expect a high level of performance to support that small number of users for that environment that doesn't require fault tolerance. **Quality of Service** could also be used, to ensure that this virtual machine receives a guaranteed level of network bandwidth if there is contention on that particular host.

**One Server Farm on Hyper-V Cluster:** By taking that same virtual machine, and moving the disks to a centralized, resilient storage repository, such as a SAN, or highly available file server, and by moving the VM's config and state to a Hyper-V cluster, the administrator can significantly raise the availability of the SharePoint single-VM farm. Whilst there is no application level resiliency, at either the web, application, or SQL layers within the VM, should the host fail, the VM will be automatically restarted on another available host in the cluster, without any administrator intervention.

**Two-Tier Farm on Two VMs:** In this example topology, the Web and Application server roles are installed on one VM and the SQL Server Database server role is installed on the second VM. This type of topology could support up to 10,000 users, but has no resiliency if these virtual machines are deployed on the same host.

Figure 36: Two Server Farm in Two VMs

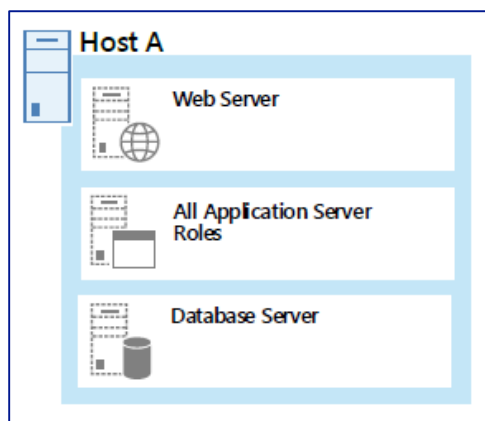


If this is your preferred deployment option, and fault tolerance is not required, aside from the capabilities described earlier to enhance performance, the administrator could also utilize **Weights and Reserves**, within the Hyper-V CPU settings for the VM, to ensure that the SQL Server VM received adequate resources to ensure high performance.

**Two-Tier Farm on Two VMs on a Hyper-V Cluster:** If the VMs are running on a cluster, this will provide a higher level of availability, however there is little use in the Anti-Affinity rules here, as there are no duplicate tiers, i.e. there is only a single web/application server, and a single SQL Server, so keeping them apart on separate hosts make little impact, as losing one, the other, or both of the VMs will essentially render the SharePoint infrastructure down. One cluster setting that may be appropriate however, would be the **Failover Priority** setting, which would help to ensure that upon failover, the SQL Server database virtual machine started first, and then the web/application server VM started afterwards.

**Three-Tier Farm on Three VMs:** In this topology, all three SharePoint server roles are placed on three different VMs on single Hyper-V host. This type of topology still lacks SharePoint resiliency, but would offer a greater level of scalability than what was offered previously when roles were combined.

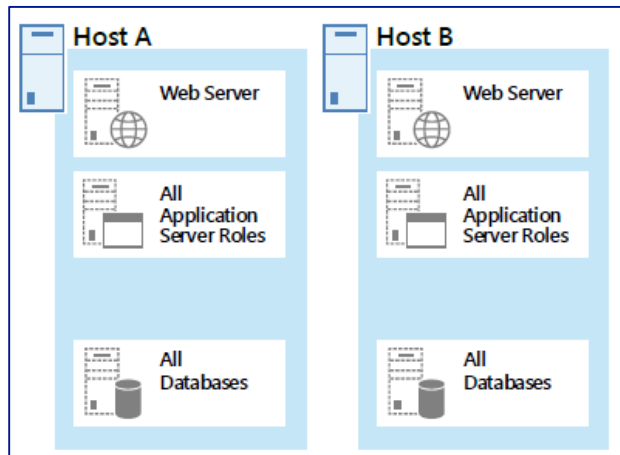
Figure 37: Three Tier Farm in 3 VMs on Single Host



**Smallest Fault Tolerant Farm using Virtualization:** This topology has a minimum of two hosts, each with a Web Server, Application Server and Database server on three different VMs. This type of topology provides fault tolerance using the minimum number of servers or hosts. SQL Server 2012 is installed and

configured on the Database server roles to support SQL clustering, mirroring, or AlwaysOn Availability Groups. Only one instance of SQL Server is deployed to each host and it is recommended that you do not deploy more than one SQL Server guest per host.

Figure 38: Smallest Fault Tolerant Farm using Virtualization



While designing availability requirements for a virtualized SharePoint 2013 farm, you need to understand the availability for each layer because this will help you to deliver a comprehensive solution that minimizes downtime. Not all solutions in an organization are likely to require the same level of availability. Required availability depends largely on the tier and number of servers and services in the SharePoint farm.

### Best Practices and Recommendations

Only one instance of SQL Server is deployed to each host. For small and medium virtual environments, it is recommended that you do not deploy more than one SQL Server guest per host to ensure that in the event of host failure, whole farms are not rendered inoperable for a period of time.

Both host servers include more memory to accommodate the number of virtual servers, including SQL Server.

In this example, the administrator could create 3 Anti-Affinity Groups:

```
$wfeAntiAffinity = New-Object System.Collections.Specialized.StringCollection
$wfeAntiAffinity.Add("SharePoint WFE")

$appAntiAffinity = New-Object System.Collections.Specialized.StringCollection
$appAntiAffinity.Add("SharePoint APP")

$dbAntiAffinity = New-Object System.Collections.Specialized.StringCollection
$dbAntiAffinity.Add("SharePoint DB")
```

Now that we have the affinity class names defined, we can assign these to the cluster groups. Once again, we can use the *Get-ClusterGroup* cmdlet to update the value of this property for each virtual machine:



```
(Get-ClusterGroup -Name SP-WFE1).AntiAffinityClassNames = $wfeAntiAffinity
(Get-ClusterGroup -Name SP-WFE2).AntiAffinityClassNames = $wfeAntiAffinity
(Get-ClusterGroup -Name SP-APP1).AntiAffinityClassNames = $appAntiAffinity
(Get-ClusterGroup -Name SP-APP2).AntiAffinityClassNames = $appAntiAffinity
(Get-ClusterGroup -Name SP-SQL1).AntiAffinityClassNames = $dbAntiAffinity
(Get-ClusterGroup -Name SP-SQL2).AntiAffinityClassNames = $dbAntiAffinity
```

In addition, customers could take advantage of the **Failover Priority** to help control the start-up ordering of the VMs upon failover.

For guidance on SQL AlwaysOn inside virtual machines, refer to the [Best Practices for Virtualizing and Managing SQL Server](#) [whitepaper](#).

This particular design, split across 2 nodes, could be scaled out with additional Web, Application, and SQL Servers, as the demand increases, or as the need for increased resiliency increases.

These example topologies are just some of the possible options that could be deployed within your environment. The Microsoft Download Center has a number of other pieces of guidance around different topologies, particularly in physical environments, which could be easily adapted and shaped for deployment into a virtualized environment. These documents include a [Streamlined Deployment Topology](#), and more of the [Traditional Deployment Topologies](#). In addition, there's a number of other architectural documents available at the [SharePoint TechCenter](#). As mentioned earlier, these topologies are just one part of the overall project for virtualizing SharePoint, with another being the significant capacity planning and testing that must take place regardless of whether deploying into a virtualized, or a physical environment. Information such as that found on [TechNet can help considerably](#) with that part of the process.

# 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 SharePoint 2013 in a virtualized environment.

## Comprehensive Management Capabilities

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 SP1, 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.

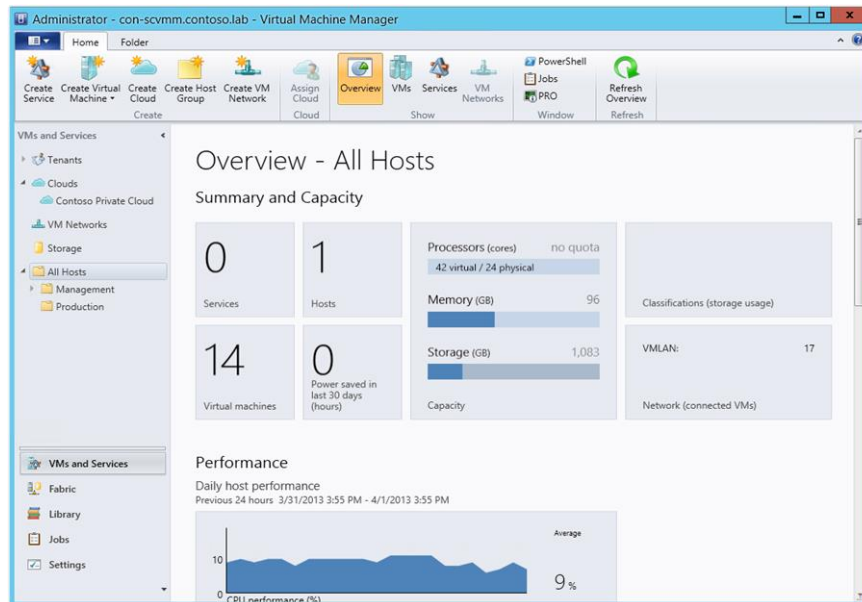
## Virtual Machine Manager

System Center 2012 SP1 Virtual Machine Manager (VMM) is the control center of a virtualized SharePoint 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.

## Centralized Fabric Configuration

As shown in Figure 39, VMM enables IT administrators to easily configure and manage virtualization hosts, networking, and storage resources in order to rapidly create and deploy SharePoint Server virtual machines and services.

Figure 39: 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 SharePoint in a virtualized environment.

### Physical-to-Virtual Conversions

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

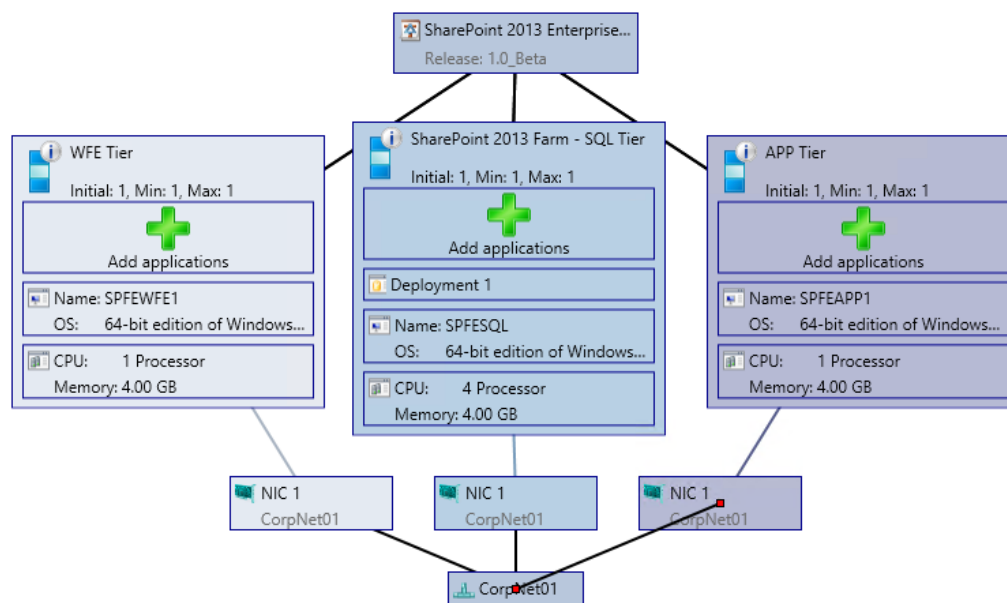
- **Online conversion:** With an online conversion, the source computer continues to perform normal operations and is available throughout the process. VMM creates a copy of local NTFS volumes and data for VSS-aware applications. VMM uses the Volume Shadow Copy Service (VSS) to ensure that data is backed 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 SharePoint 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 SharePoint Server environments.

### Virtual Machine Profiles and Templates

To streamline deployment, VMM offers profiles, virtual machine templates, and service templates:

- **Profile:** A profile is a library resource containing specifications that can be applied to a new virtual machine or virtual machine template. Examples could be a Guest OS profile, which contains items such as domain join information, or a product key for activation, or a hardware profile, which contains a specific hardware configuration with a certain number of vCPUs, disks etc.
- **Virtual machine template:** Virtual machine templates encapsulate a standard set of configuration settings that can be used when creating a virtual machine. These templates can help IT administrators to quickly create virtual machines with consistent hardware and operating system settings. This can be extremely useful for the rapid deployment of SharePoint virtual machines into an infrastructure. Virtual machine templates can also be used to restrict the settings available to self-service users creating new virtual machines.
- **Service template:** The service template for SharePoint 2013 Enterprise provides a three-tier architecture that can be used to simplify the deployment of SharePoint Server services in a private cloud environment<sup>52</sup> (Figure 40). Service templates are an integral part of System Center 2012: They are used to automate tasks and reduce deployment time and cost by dynamically provisioning services. These templates also can be extended to automate more advanced deployment scenarios, as required. Service templates can be created using the Service Template Designer in VMM. Once a service template is generated, SharePoint 2013 tiers and networking components can be added.

Figure 40: Service template for SharePoint 2013 Enterprise three-tier architecture



The VMM Service Template example for SharePoint 2013 can be [downloaded from the TechNet gallery](#), and the [accompanying blog post](#) can help guide through the installation and configuration.

## Virtual Machine Deployment

A virtual machine template can accelerate deployment of SharePoint 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.

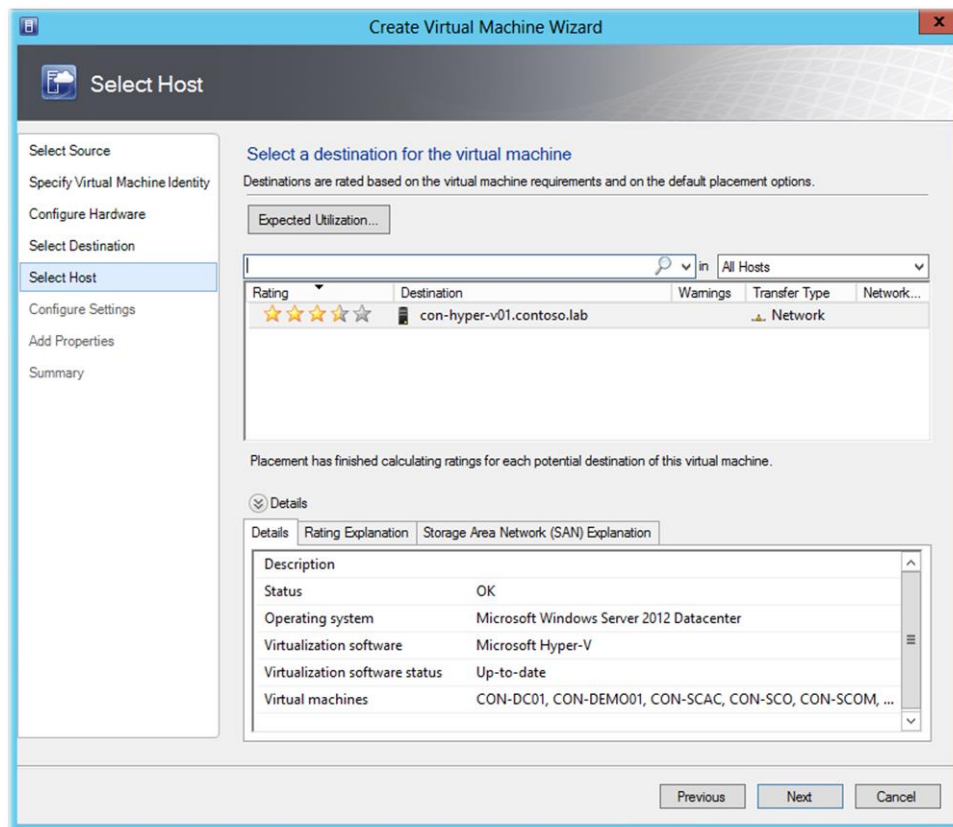
### Intelligent Placement

VMM can help to identify the most appropriate physical host servers for virtualized SharePoint Server workloads (Figure 41). 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.

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.

Figure 41: Intelligent Placement in VMM



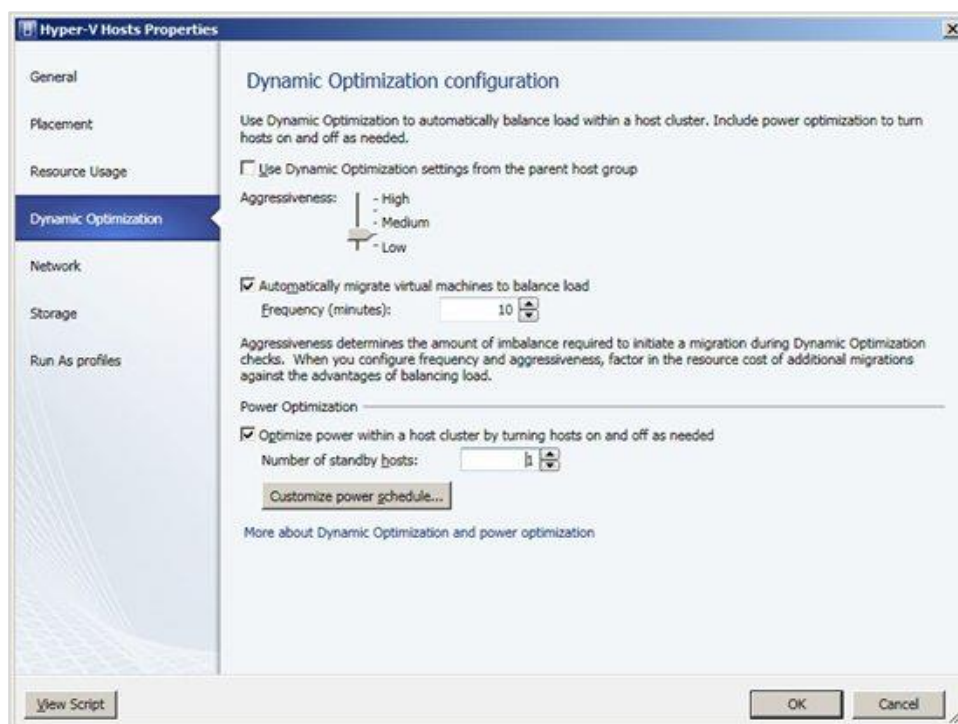
## Storage Classification

VMM also provides the ability for an administrator to apply simple classifications to storage, which can be used for storing and running SharePoint 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 virtual machine templates so that VMM automatically ensures that the chosen type of storage will be used for a particular deployment.

## Dynamic Optimization

Once SharePoint 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 42). For example, you may have several hosts in a cluster, but one of the hosts has some SharePoint 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 virtual machines. This helps to ensure that workloads inside the virtual machines always receive the resources they need to meet demand, without impacting other workloads running on the cluster.

Figure 42: 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, be sure to 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.

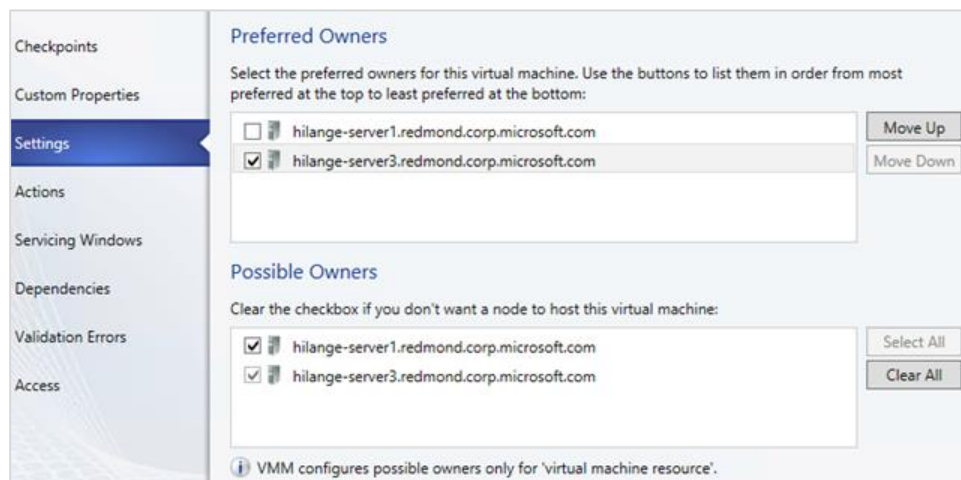
## Virtual Machine Priority and Affinity

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



VMM also provides the ability for an 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 43). 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.

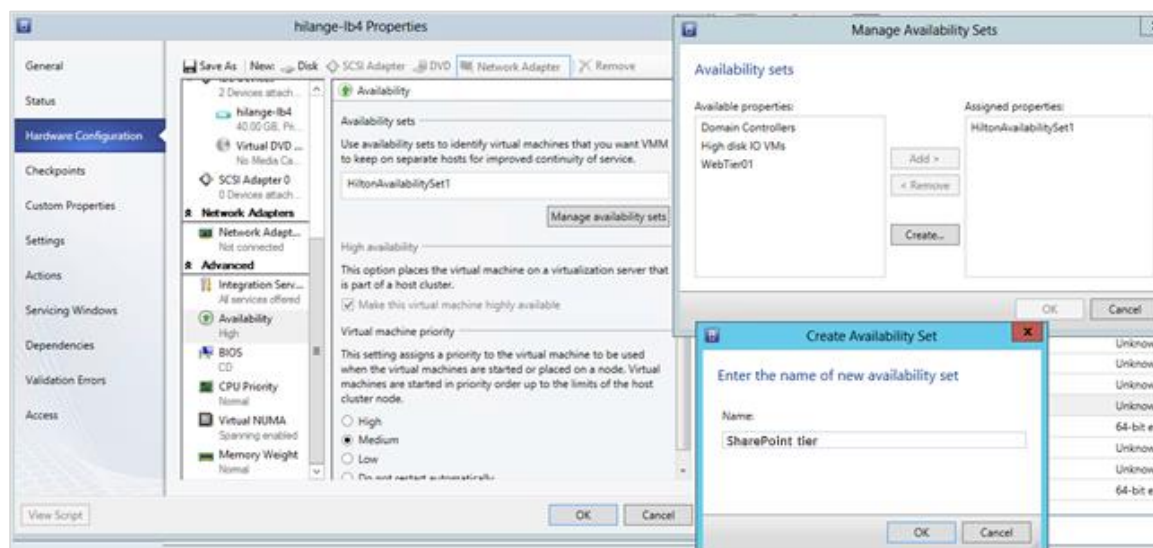
Figure 43: Defining preferred and possible owners in VMM



## Availability Sets

When an IT administrator places multiple virtual machines in an availability set, VMM attempts to keep those virtual machines on separate hosts whenever possible (Figure 43). 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 as we learnt earlier. 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 47: Availability set for SharePoint virtual machines





## Best Practices and Recommendations

When creating a guest cluster (such as a SQL Server AlwaysOn Availability Group for SharePoint databases) a set of Web Front Ends, or a set of Application Servers, virtualized on top of a Hyper-V host cluster, it can be beneficial to keep the individual SharePoint Server VMs on separate hosts. If one physical host is lost, it takes down only a single VM of the Web Front Ends, Application Servers, or AlwaysOn Availability Group because the availability set within VMM ensures that related VMs are running on separate hosts in the Hyper-V cluster.

## Private 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 IT administrator can quickly and easily manage the private cloud definition, access to the private cloud itself, and the underlying physical resources (Figure 48). VMM also provides granular, role-based access to end users, application owners, or SharePoint administrators.

Figure 48: 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.
- **Elasticity:** Administrators can add resources to a private cloud to increase capacity.
- **Optimization:** Use of underlying resources is continually optimized without affecting the overall private cloud user experience.

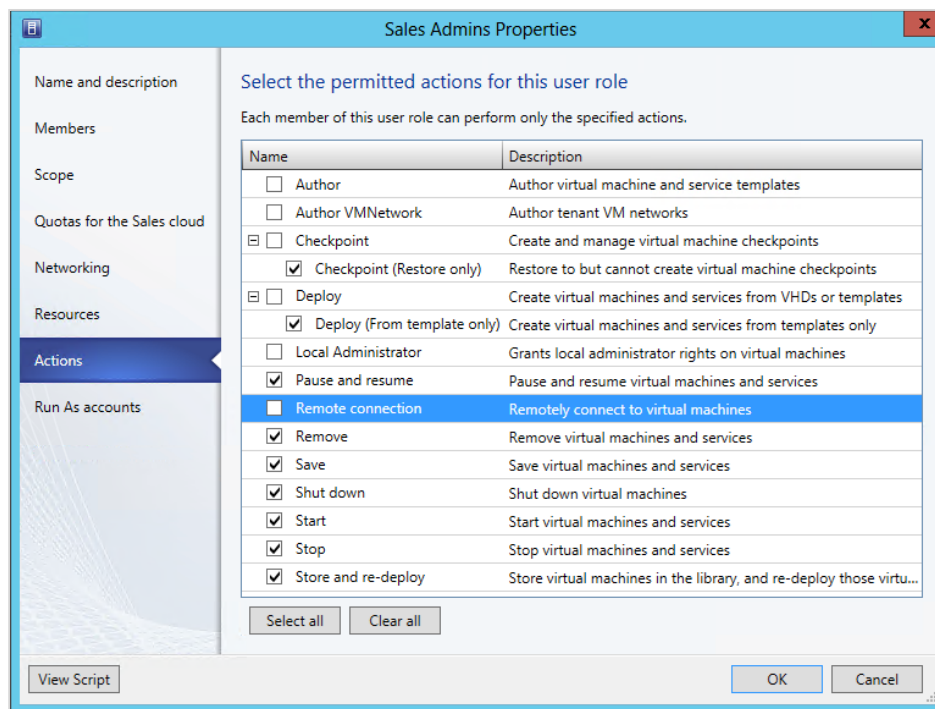
When creating a private cloud, you select the underlying fabric resources that will be available, configure library paths for users, and set the capacity. Therefore, before you create a private cloud, you should configure the fabric resources, such as storage, networking, library servers and shares,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

## Best Practices and Recommendations

From a SharePoint perspective, an IT administrator can define a cloud that can be used exclusively with SharePoint 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 this cloud. This ensures that the only virtual machines that can be deployed into the cloud are those that already have SharePoint within the templates, thereby optimizing deployment.

Once the private cloud is created, an IT administrator can assign access to certain users and groups, such as SharePoint administrators, within the IT infrastructure. Through rich, granular role-based controls, the IT administrator can delegate who can see what inside the cloud, as well as who can perform which tasks associated with it (Figure 44). 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 44: Configuring user role actions in VMM

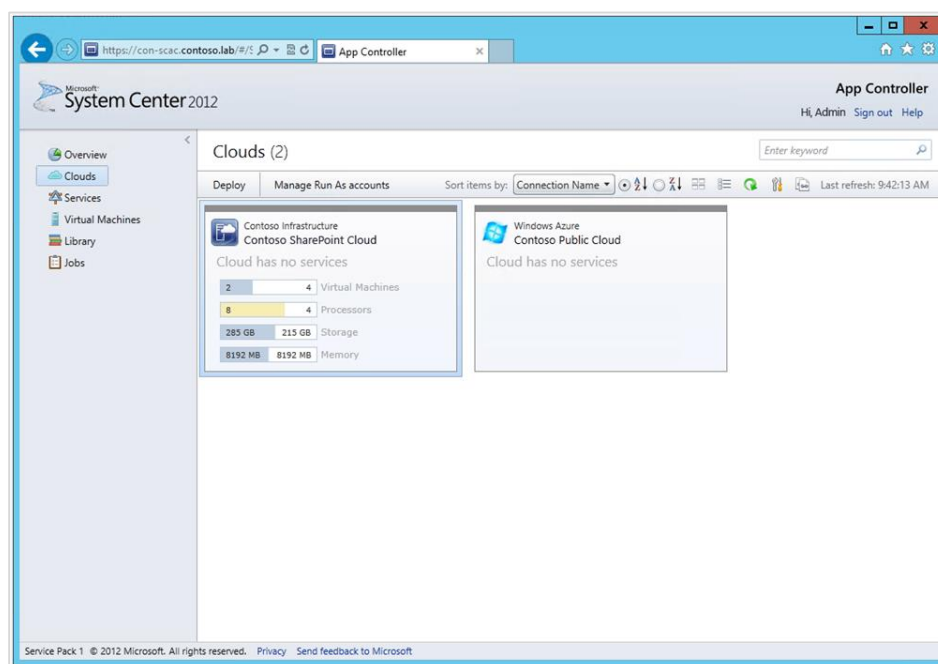


# App Controller

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 SharePoint 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.

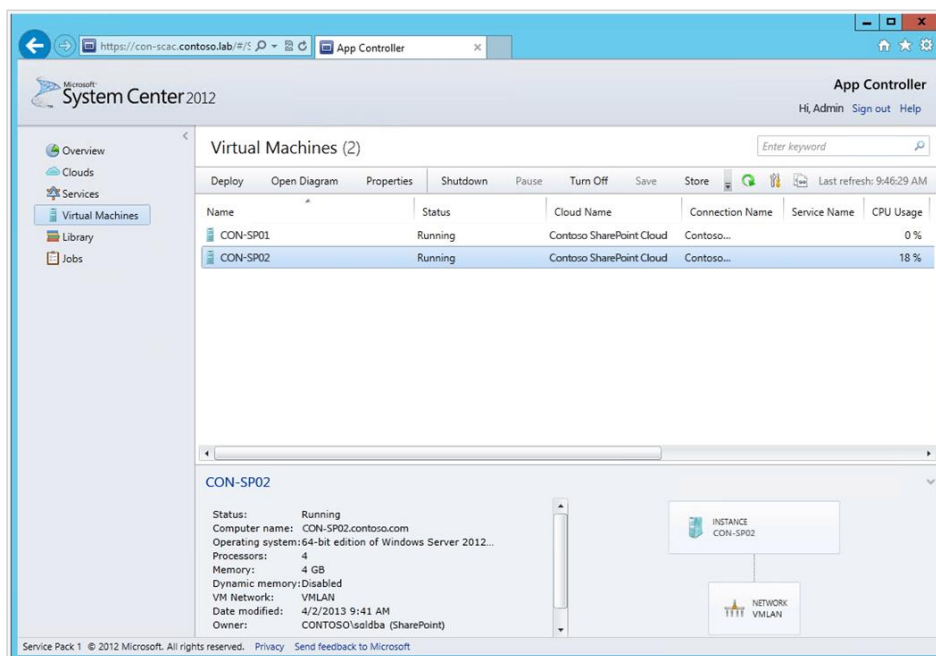
The following example shows the steps that a SharePoint administrator might take to self-provision a SharePoint Server virtual machine using App Controller. When the administrator first logs on to the App Controller interface, an overview screen is presented (Figure 45). This screen is dynamically generated based on identity; therefore, it automatically displays what can be accessed. By selecting the *Clouds* option on the left side of the screen, the administrator has a visible representation of accessible clouds. In this example, note that the administrator has access not only to the Contoso SharePoint Cloud, but also to the Contoso Public Cloud – a cloud that’s running in Windows Azure, in a Microsoft datacenter.

Figure 45: Using App Controller to access a SharePoint private cloud



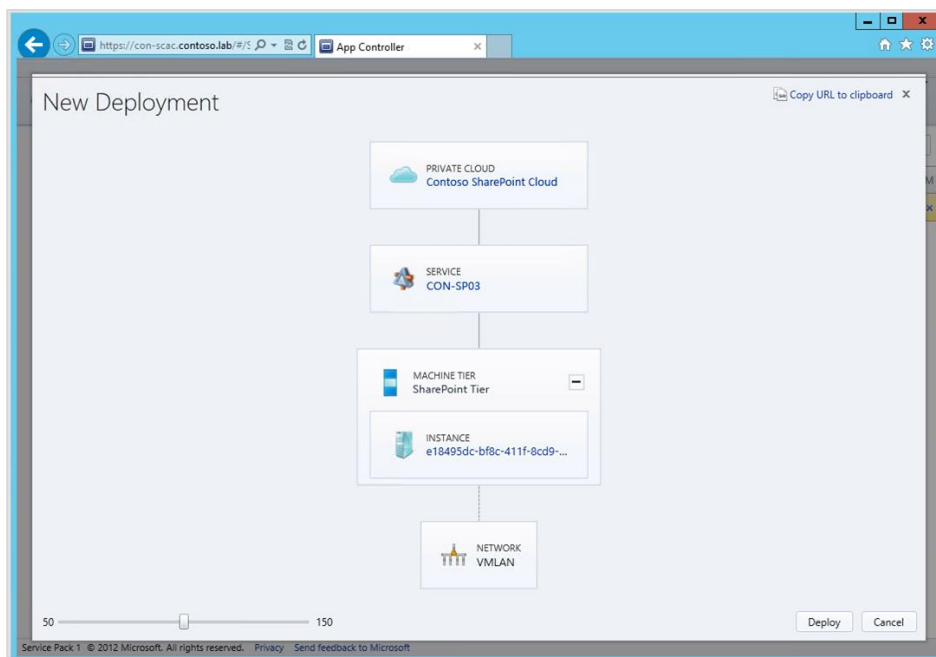
In Figure 46, the SharePoint administrator selects *Virtual Machines* on the left side of the screen. This brings up the list of current virtual machines. What is important to note here is that the SharePoint administrator only sees virtual machines that IT administrators have specifically provided and 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 SharePoint administrator is only able to perform certain tasks on the virtual machines. In this example, the tasks for start, stop, shut down, and deploy are enabled, but not pause or save.

Figure 46: Virtual Machines view in App Controller



By selecting a particular cloud, the SharePoint 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. When the SharePoint administrator clicks *Deploy* to start the virtual machine provisioning process, VMM automatically orchestrates the deployment and placement of the SharePoint Server virtual machine (Figure 47). Once the new virtual machine is deployed, the SharePoint administrator can access it through App Controller and perform the tasks and actions that the IT administrators have enabled.

Figure 52: Deploying a SharePoint Server virtual machine with App Controller



# Service Delivery and Automation

To review, the subsections above have discussed:

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

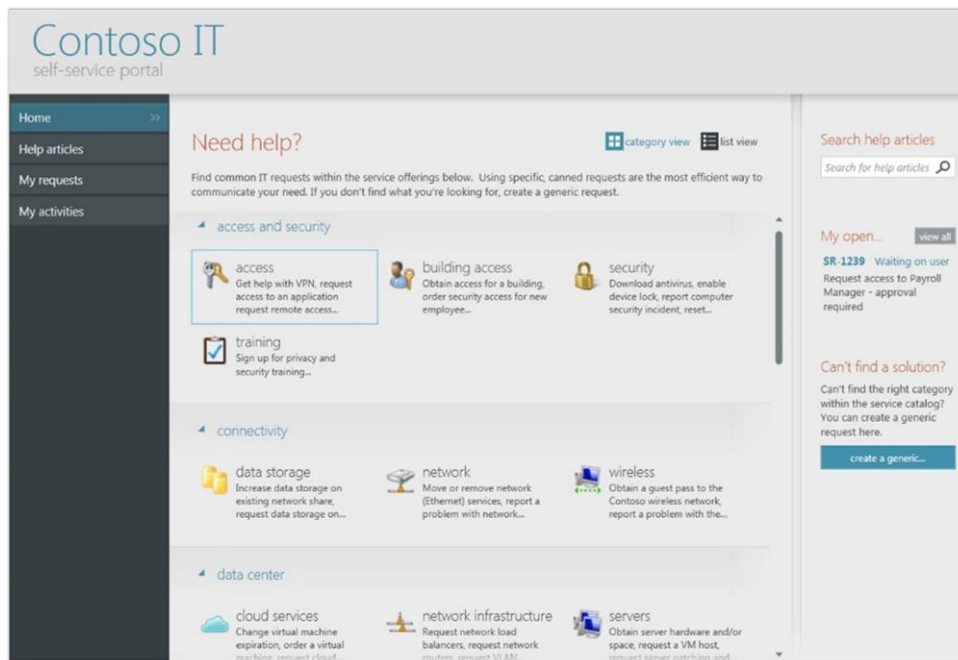
Now consider the scenario where IT administrators want to use App Controller to allow SharePoint administrators to access their virtual machines, but they also want to enact a mechanism through which the SharePoint administrators must request new virtual machines, as needed, instead of creating them at will. To manage this scenario, your organization needs the Service Manager and Orchestrator components of System Center, as well as the no-cost 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 SharePoint administrators.

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.
- **IT as a service:** Service Manager includes a rich self-service portal that provides role-based access to the service catalog (Figure 48). 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 48: Self-service portal in Service Manager

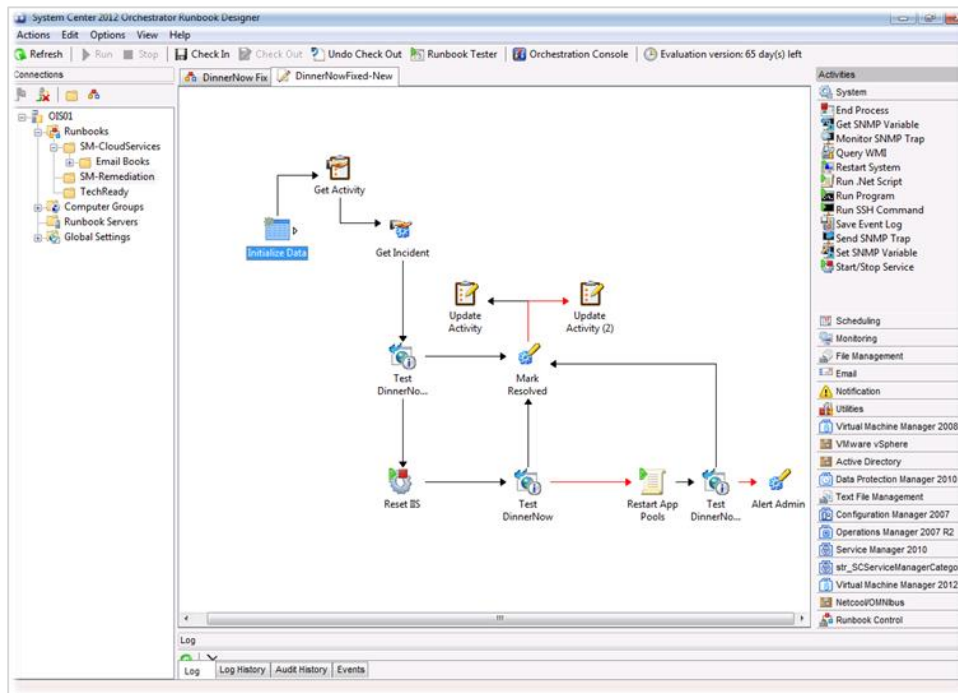


- **Integration:** Connectors simplify and streamline integration between Service Manager and other System Center components. You can use Service Manager connectors to import data as configuration items from Active Directory Domain Services, Configuration Manager, Orchestrator, VMM, and Operations Manager. 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 49). 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 49: 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 more 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 SharePoint 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 SharePoint administrators who want to request infrastructure to run their workloads. In Figure 50, when the SharePoint administrator 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 50: Self-service portal in Service Manager

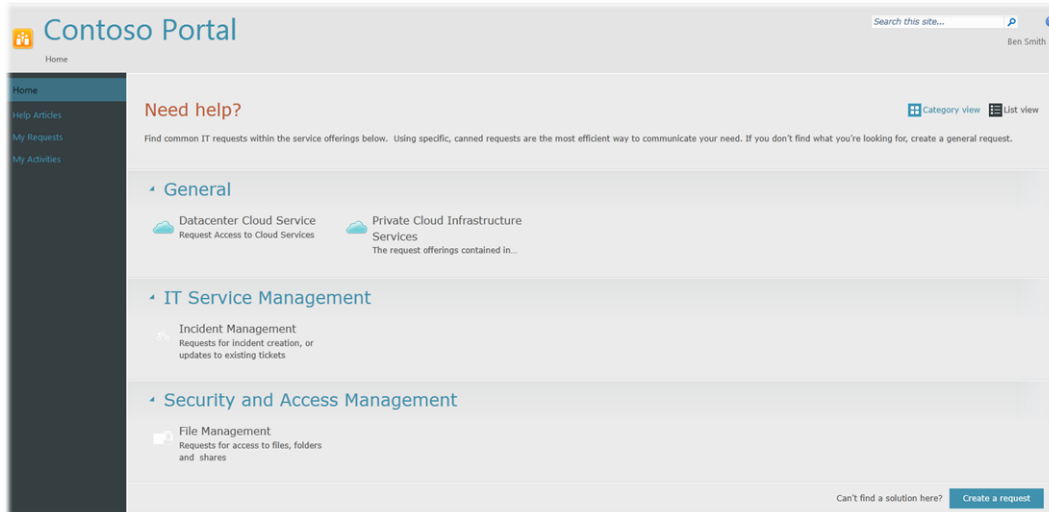
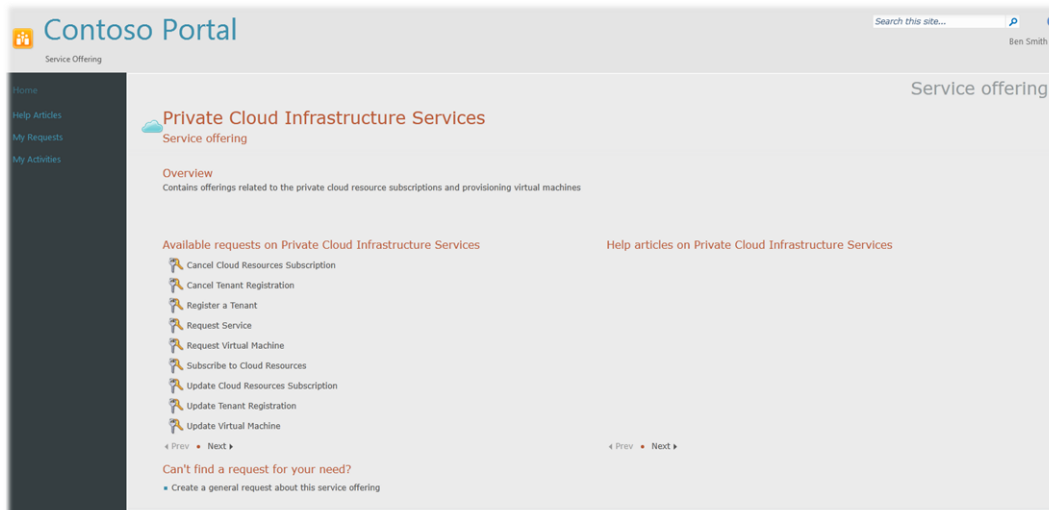


Figure 56 shows the *Service Offering* page of the portal. Service Offerings essentially group together requests that the specific user can make. In this example, the SharePoint administrator selects the Service Offering entitled *Private Cloud Infrastructure Services* to be presented with available requests.

Figure 56: Service Offerings page and related requests in Service Manager



The available requests are essentially the menu of choices that IT has provided for the SharePoint administrator. 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 SharePoint perspective, IT can define specific requests that relate to SharePoint. These requests can be general or specific—for example, a request for IT to create a pool of resources for the SharePoint team, or a request to deploy a specific virtual machine. Remember the [previous example](#) where the SharePoint administrator used App Controller to deploy a virtual machine from a service template. In the current example, the same process is taking place, but the SharePoint administrator is making a request to have the activity performed (Figure 51). If the request is approved, Orchestrator works in conjunction with Service Manager and VMM to create and deliver the virtual machine. From there, the SharePoint administrator can interact with the virtual machine through App Controller.

Figure 51: Making a cloud resource request in Service Manager

**Service Request**

**Production Cloud Resource Request**

This request belongs to: [Private Cloud Services](#)

Please complete the following in full

Select the production cloud you want access to

Select the cloud capacity you need

Select Recovery Service Level

Antivirus Protection and Application Monitoring Automatically included

1 **Provide information**

2 Review and submit

3 Confirmation

← Back   Next →   Cancel ✕

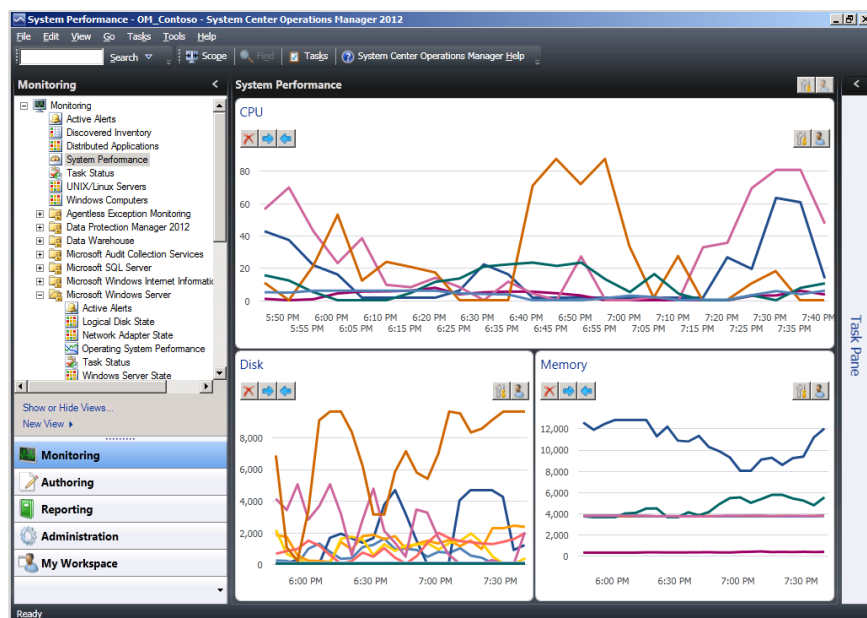
Note that the SharePoint administrator can request a private cloud and specify its capacity, recovery service level, and antivirus and monitoring status. Once the request is submitted, the appropriate IT administrator is notified to initiate approval. Upon approval, Orchestrator, as part of the integrated 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 SharePoint administrator 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 52). 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 52: Dashboard in Operations Manager



## SharePoint Server 2013 Management Pack

The ability to provide end-to-end management across infrastructure is a critical step in ensuring the health of hosts and clusters, virtual machines, and the private cloud itself. Further, Operations Manager can be used to monitor mission-critical SharePoint Server workloads using the System Center Management Pack for SharePoint 2013. This management pack provides the capabilities for Operations Manager 2007 R2 and Operations Manager 2012 to monitor SharePoint 2013 products, including SharePoint Server, Project Server, and Search Server.

The SharePoint Server 2013 Management Pack is designed to be used for monitoring SharePoint 2013 Products events, collecting SharePoint component-specific performance counters in one central location, and for raising alerts for operator intervention as necessary. By detecting, sending alerts, and

automatically correlating critical events, this management pack helps indicate, correct, and prevent possible service outages or configuration problems, allowing you to proactively manage SharePoint servers and identify issues before they become critical. The management pack monitors and provides alerts for automatic notification of events indicating service outages, performance degradation, and health monitoring.

- Health monitoring of SharePoint Server 2013
- Monitors Events and Services and alerts when service outages are detected
- Monitors Performance and warns users when SharePoint performance is at risk
- Forwards users to up-to-date TechNet knowledge articles

The SharePoint Server 2013 Management Pack is composed of three elements: discovery, monitoring, and rules:

- **Discovery:** The management pack discovers service applications and features in SharePoint farms.
- **Monitoring:** The management pack monitors SharePoint services and features to create incidents that can be used for health monitoring and performance. It also forwards up-to-date TechNet articles for known incidents.
- **Rules:** Rules collect data when predefined events occur and action is taken. For example, a rule for Microsoft Access Data Services collects data when a web-based front-end machine is unable to communicate with a specific back-end Access Data Services application server. To correct the issue, traffic is load balanced to another server, if one is available.

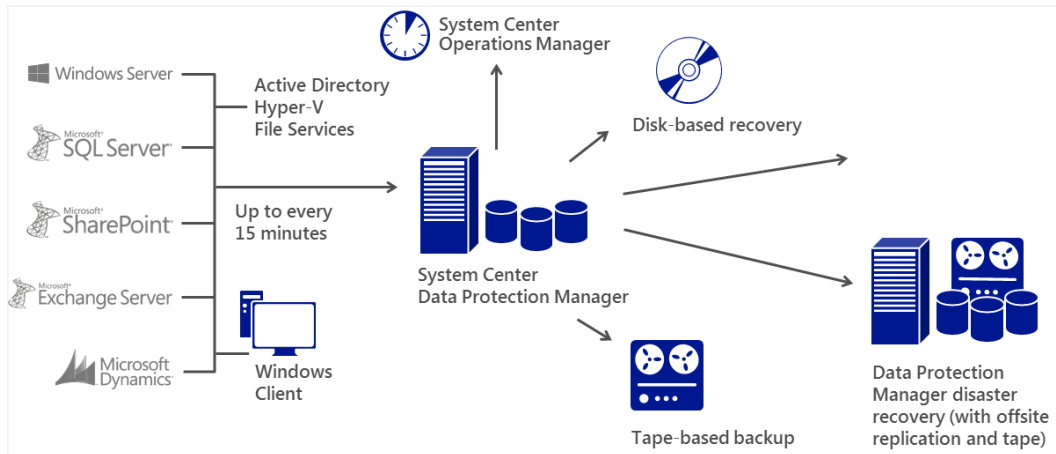
## Data Protection Manager

System Center 2012 SP1 Data Protection Manager (DPM) provides comprehensive disk and tape-based data protection and recovery across Active Directory domains for SharePoint virtual infrastructure. DPM continuously performs synchronization and creates multiple recovery points to more reliably and rapidly restore mission-critical data at any time. DPM also can centrally manage system state and bare metal recovery (BMR), and provides snapshot-based backup of SharePoint farm content by using VSS within the Windows Server operating system.

As shown in Figure 53, DPM provides the following:<sup>53</sup>

- Quick and frequent backup of the entire SharePoint farm
- Capabilities for item-level recovery
- Capabilities for backup and recovery from disk
- Granular retention policies

Figure 53: Key capabilities of Data Protection Manager



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. DPM natively protects the following key components of a SharePoint farm:

- SharePoint farm content including automatic protection of new content DBs
- Front-end web server content
- SharePoint Search Index

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.

# Conclusion

With an effective virtualization strategy, organizations can consolidate SharePoint 2013, along with other key workloads on Windows Server 2012 Hyper-V to help remediate server sprawl, maximize ROI, and decrease TCO and operational costs. It is essential to properly plan the virtualization of mission-critical workloads like SharePoint, and the best practices and recommendations detailed in this guide can help with such planning. The fabric considerations focus on planning the physical infrastructure—including processor, memory, storage, and network—while the SharePoint VM, resiliency and topology considerations describe how to configure virtual machines using SharePoint 2013 and Windows Server 2012 Hyper-V.

System Center 2012 SP1 should be used to proactively manage, monitor, and protect the SharePoint environment. System Center 2012 SP1 provides a powerful set of tools to manage virtual resources. The purpose of these tools is comprehensive, ranging from assisting with P2V conversions, to rapidly provisioning virtual machines and services from templates, and protecting key resources within the environment.

Together, SharePoint 2013, Windows Server 2012, and System Center 2012 SP1 provide an integrated information platform that spans 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 SharePoint 2013**

<http://sharepoint.microsoft.com/en-in/Pages/default.aspx>

**Microsoft SharePoint 2013 TechNet**

<http://technet.microsoft.com/en-us/sharepoint/fp142366>

**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

- <sup>1</sup> ESG Lab. "Virtualizing SharePoint Workloads with Microsoft Hyper-V R2."  
[http://download.microsoft.com/download/1/7/F/17FB551C-0905-4A04-AB46-2EBA616CFDF3/ESG%20Preso%20Microsoft%20Hyper-V%20Performance%20SharePoint%20Mar%2011\\_Wide.pdf](http://download.microsoft.com/download/1/7/F/17FB551C-0905-4A04-AB46-2EBA616CFDF3/ESG%20Preso%20Microsoft%20Hyper-V%20Performance%20SharePoint%20Mar%2011_Wide.pdf)
- <sup>2</sup> Microsoft TechNet. "Capacity Management and Sizing Overview for SharePoint Server 2013." July 2012.  
<http://technet.microsoft.com/en-us/library/ff758647.aspx>
- <sup>3</sup> ESG Lab. "Virtualizing SharePoint Workloads with Microsoft Hyper-V R2."  
[http://download.microsoft.com/download/1/7/F/17FB551C-0905-4A04-AB46-2EBA616CFDF3/ESG%20Preso%20Microsoft%20Hyper-V%20Performance%20SharePoint%20Mar%2011\\_Wide.pdf](http://download.microsoft.com/download/1/7/F/17FB551C-0905-4A04-AB46-2EBA616CFDF3/ESG%20Preso%20Microsoft%20Hyper-V%20Performance%20SharePoint%20Mar%2011_Wide.pdf)
- <sup>4</sup> Microsoft. "SharePoint 2010 Virtualization Guidance and Recommendations."  
<http://download.microsoft.com/download/D/D/E/DDED76C5-9874-473F-99C1-B902D325ED38/bi/oit2010-whitepaper-virtualization-guidance.pdf>
- <sup>5</sup> Gartner. "Magic Quadrant for x86 Server Virtualization Infrastructure." Jun 2012.  
<http://www.gartner.com/technology/reprints.do?id=1-1AVRXJO&ct=120612&st=sb>
- <sup>6</sup> Microsoft. "System Center 2012 - Infrastructure Management."  
<http://www.microsoft.com/en-us/server-cloud/system-center/infrastructure-management.aspx>
- <sup>7</sup> 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>
- <sup>8</sup> Microsoft. "Installing Windows Server 2012." May 2012.  
<http://technet.microsoft.com/en-us/library/ji134246.aspx>
- <sup>9</sup> Microsoft TechNet. "Windows Server Installation Options."  
<http://technet.microsoft.com/en-us/library/hh831786.aspx>
- <sup>10</sup> Microsoft MSDN. "Minimal Server Interface for Windows Server 2012."  
[http://msdn.microsoft.com/en-us/library/windows/desktop/hh846317\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/hh846317(v=vs.85).aspx)
- <sup>11</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>12</sup> 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>
- <sup>13</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>14</sup> Microsoft TechNet. "Configure the Processors for the Virtual Machines."  
<http://technet.microsoft.com/en-us/library/ff621103.aspx#CPUConfig>
- <sup>15</sup> 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>
- <sup>16</sup> 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](http://blogs.msdn.com/b/virtual_pc_guy/archive/2011/02/14/hyper-v-cpu-scheduling-part-1.aspx)
- <sup>17</sup> Microsoft MSDN. "Understanding Non-Uniform Memory Access."  
<http://msdn.microsoft.com/en-us/library/ms178144.aspx>
- <sup>18</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>19</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)



- 
- <sup>20</sup> 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>
- <sup>21</sup> ten Seldam, Matthijs (TechNet Blog). "Windows Server – Sockets, Logical Processors, Symmetric Multi-Threading." Oct 2012.  
<http://blogs.technet.com/b/matthijs/archive/2012/10/14/windows-server-sockets-logical-processors-symmetric-multi-threading.aspx>
- <sup>22</sup> Microsoft TechNet. "Storage Spaces – Designing for Performance."  
<http://social.technet.microsoft.com/wiki/contents/articles/15200.storage-spaces-designing-for-performance.aspx>
- <sup>23</sup> Microsoft TechNet. "Storage Spaces – Designing for Performance."  
<http://social.technet.microsoft.com/wiki/contents/articles/15200.storage-spaces-designing-for-performance.aspx>
- <sup>24</sup> 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>
- <sup>25</sup> Microsoft TechNet. "Failover Clustering Hardware Requirements and Storage Options." Aug 2012.  
<http://technet.microsoft.com/en-in/library/jj612869.aspx>
- <sup>26</sup> Microsoft. "Networking – Windows Server 2012."  
[http://download.microsoft.com/download/7/E/6/7E63DE77-EBA9-4F2E-81D3-9FC328CD93C4/WS%202012%20White%20Paper\\_Networking.pdf](http://download.microsoft.com/download/7/E/6/7E63DE77-EBA9-4F2E-81D3-9FC328CD93C4/WS%202012%20White%20Paper_Networking.pdf)
- <sup>27</sup> 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](http://download.microsoft.com/download/B/3/0/B301FE04-ABAB-4C89-8631-0FCBE4147B84/WS_2012_Product_White_Paper.pdf)
- <sup>28</sup> Microsoft TechNet. "Cluster-Aware Updating Overview."  
<http://technet.microsoft.com/en-us/library/hh831694.aspx>
- <sup>29</sup> 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>
- <sup>30</sup> Microsoft TechNet. "What's New in Failover Clustering." Sep 2012.  
[http://technet.microsoft.com/en-us/library/hh831414.aspx#BKMK\\_PLACE](http://technet.microsoft.com/en-us/library/hh831414.aspx#BKMK_PLACE)
- <sup>31</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>32</sup> 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>
- <sup>33</sup> 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-a026-5bfcd076d9b9/sql2008inhyperv2008.docx>
- <sup>34</sup> McMurray, Robert. "IIS 8.0 Multicore Scaling on NUMA Hardware." Feb 2012.  
<http://www.iis.net/learn/get-started/whats-new-in-iis-8/iis-80-multicore-scaling-on-numa-hardware>
- <sup>35</sup> Posey, Brien (TechNet Blog). "Virtualization: Optimizing Hyper-V Memory Usage."  
<http://technet.microsoft.com/en-us/magazine/hh750394.aspx>
- <sup>36</sup> 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](http://download.microsoft.com/download/E/0/5/E05DF049-8220-4AEE-818B-786ADD9B434E/Implementing_and_Configuring_Dynamic_Memory.docx)
- <sup>37</sup> Microsoft TechNet. "Hyper-V Dynamic Memory Overview." Feb 2012.  
<http://technet.microsoft.com/en-us/library/hh831766.aspx>
- <sup>38</sup> 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>
- <sup>39</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>40</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>41</sup> Microsoft TechNet. "Configure the Memory for the Virtual Machines." Jan 2013.  
<http://technet.microsoft.com/en-us/library/ff621103.aspx#ConfigMem>

- 
- <sup>42</sup> Microsoft. "Running SQL Server with Hyper-V Dynamic Memory." Jul 2011.  
<http://download.microsoft.com/download/D/2/0/D20E1C5F-72EA-4505-9F26-FF9550EFD44/Best%20Practices%20for%20Running%20SQL%20Server%20with%20HVDM.docx>
- <sup>43</sup> 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>
- <sup>44</sup> Microsoft. "SharePoint 2010 Virtualization Guidance and Recommendations."  
<http://www.microsoft.com/en-us/download/details.aspx?id=19508>
- <sup>45</sup> Microsoft TechNet. "Use best practice configurations for the SharePoint 2013 virtual machines and Hyper-V environment." Jul 2012.  
<http://technet.microsoft.com/en-us/library/ff621103.aspx#ConfigNet>
- <sup>46</sup> 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-io-v-in-Hyper-V-part-5.aspx>
- <sup>47</sup> Howard, John (Microsoft TechNet Blog). "Everything you wanted to know about SR-IOV in Hyper-V Part 6." Mar 2012.  
<http://blogs.technet.com/b/jhoward/archive/2012/03/19/everything-you-wanted-to-know-about-sr-io-v-in-hyper-v-part-6.aspx>
- <sup>48</sup> ten Seldam, Matthijs (Microsoft TechNet Blog). "Windows Server 2012 Hyper-V and SR-IOV."  
<http://blogs.technet.com/b/matthijs/archive/2012/10/13/windows-server-2012-hyper-v-and-sr-io-v.aspx>
- <sup>49</sup> Microsoft MSDN. "SR-IOV VF Failover and Live Migration Support." May 2013.  
[http://msdn.microsoft.com/en-in/library/windows/hardware/hh440249\(v=vs.85\).aspx](http://msdn.microsoft.com/en-in/library/windows/hardware/hh440249(v=vs.85).aspx)
- <sup>50</sup> Microsoft. "Server Virtualization – Windows Server 2012."  
[http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper\\_Hyper-V.pdf](http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf)
- <sup>51</sup> Microsoft TechNet. "Quality of Service (QoS) Overview." Mar 2013.  
[http://technet.microsoft.com/en-us/library/hh831679.aspx#bkmk\\_bandwidth](http://technet.microsoft.com/en-us/library/hh831679.aspx#bkmk_bandwidth)
- <sup>52</sup> Microsoft TechNet Blog. "Application Management-Example-Deploying a Service to Your Private Cloud (Part 4)". Apr 2013.  
<http://blogs.technet.com/b/privatecloud/archive/2013/04/30/application-management-example-deploying-a-service-to-your-private-cloud-part-4.aspx>
- <sup>53</sup> Microsoft TechNet. "What's New in System Center 2012 - Data Protection Manager." Jan 2013.  
<http://technet.microsoft.com/en-us/library/hh848299.aspx>