Microsoft | Hyper-V Cloud

HYPER-V CLOUD DEPLOYMENT GUIDES

MODULE 1: ARCHITECTURE AND SIZING

The Hyper-V Cloud Deployment Guides from Microsoft can help IT Pros by:

- *Accelerating deployment by providing best practices for planning and delivering a Private Cloud solution based on Microsoft Virtualization technologies.*
- *Reducing training costs by providing methodologies for delivering Server Virtualization scenarios.*
- *Lowering risk by providing real-world examples of problems and solutions encountered by Microsoft architects and consultants.*

INTRODUCTION

The Architecture and Sizing Guide outlines the design, hardware, software, and support considerations that must be taken into account when designing the server architecture for a private cloud infrastructure.

This guide outlines the minimum system requirements, supported operating systems, host server architecture patterns, storage design, and server design considerations for deploying **Microsoft Windows Server® 2008 R2 operating system with Hyper-V™ virtualization technologies, System Center Virtual Machine Manager 2008 R2 and System Center Virtual Machine Manager Self Service Portal 2.0.**

This Deployment Guide is one of five modules that are a part of the Microsoft Hyper-V Cloud Deployment Guides that are based on the framework that Microsoft Consulting Services has leveraged to deliver Server Virtualization for several years in over 82 countries.

CONTENTS

OVERVIEW OF COMPONENTS

Microsoft Windows Server ® 2008 R2 with Hyper-V™

The host servers are one of the critical components of a dynamic, virtual infrastructure. The host servers, running Windows Server® 2008 R2 with Hyper-V™ technology, provide the foundation for running virtual machine guests and also provide the management interface between the guests and Microsoft® System Center Virtual Machine Manager.

A detailed host server design and sizing methodology is described and a set of reference server architectures is presented. The reference server architectures are intended to be a starting point for the design process and provide a foundation for documenting the final design.

For detailed guidance on how to get started installing and configuring Microsoft Windows Server 2008 R2 Hyper-V please go to:

http://technet.microsoft.com/en-us/library/cc732470(WS.10).aspx

System Center Virtual Machine Manager 2008 R2

The primary tool for managing the virtual infrastructure will be System Center Virtual Machine Manager. System Center Virtual Machine Manager can scale across a wide range of virtual environments, ranging from a single server for smaller environments to a fully distributed enterprise environment that manages hundreds of hosts running thousands of virtual machines.

Virtual Machine Manager delivers the following key features:

- Designed for managing virtual machines running on Windows Server® 2008 Hyper–V™ and Microsoft Hyper-V Server.
- Virtualization support for virtual machines running on Microsoft Virtual Server and VMware ESX.
- End-to-end support for consolidating physical servers onto a virtual infrastructure.
- Performance and Resource Optimization (PRO) for dynamic and responsive management of virtual infrastructure (requires System Center Operations Manager).
- Intelligent Placement of virtual workloads on the best-suited physical host servers.

- A complete library to centrally manage all the building blocks of the virtual data center.

For detailed guidance on how to get started installing and configuring System Center Virtual Machine Manager 2008 R2 please go to:
http://technet.microsoft.com/en-us/systemcenter/vmm/default.aspx

SCVMM 2008 R2 Self-Service Portal 2.0

Using the Microsoft System Center Virtual Machine Manager Self-Service Portal 2.0, enterprise datacenters can provide infrastructure-as- a-Service to business units within the enterprise. The self-service portal provides a way for groups within an organization to manage their own IT needs while the centralized infrastructure organization manages a pool of physical resources (servers, networks, and related hardware).

The self-service portal has four components:

- **VMSSP website**. A Web-based component that provides a user interface to the self-service portal. Through the VMMSSP website, infrastructure administrators can perform various tasks such as pooling infrastructure assets in the self-service portal, extending virtual machine actions, creating business unit and infrastructure requests, validating and approving requests, and provisioning virtual machines (using the self-service virtual machine provisioning feature). Administrators can also use the VMMSSP website to view information related to these tasks.
- **VMMSSP database.** A SQL Server database that stores information about configured assets, information related to business units and requests, and information about what has been provisioned to various business units. The database also stores the XML that encodes default and customized virtual machine actions and other information related to the configuration of the self-service portal.
- **VMMSSP server.** A Windows service that runs default and customized virtual machine actions that the user requests through the VMMSSP website.
- **Reporting Dashboard.** A reporting service built on Windows SharePoint Services 3.0 SP2. The Dashboard provides 'out-of-box' reports and the ability to quickly produce custom reports.

Business units that enroll in the self-service portal system can use the portal to do the following:

- **Use standardized forms to request new infrastructures or changes to infrastructure components.** Each business unit can submit requests to the infrastructure administrator. The standardized forms ensure that the infrastructure administrator has all of the information needed to fulfill the

requests without needing to repeatedly contact the business unit for details.

- **Create and manage virtual machines.** The VMMSSP website includes self-service provisioning forms that business units can use to create virtual machines. When a business unit submits a request to create virtual machines, the self-service portal starts an automated provisioning process creates the virtual machines more quickly and efficiently than a manual process.
- **Delegate the details of virtual machine management.** Each business unit can designate its own administrators, advanced operators, and users.

Infrastructure administrators can use the self-service portal to do the following:

- **Extend the default virtual machine actions to fit your datacenter.** You can work with technology partners and hardware vendors to modify the standard "actions" that the self-service portal uses to create and manage virtual machines. In this way, you can extend the self-service portal to use specific storage area networks (SANs), load balancers, and so forth.
- **Simplify the process of enrolling business units and defining their needs.** The self-service portal collects information about a business unit and about the resources they want to set up.
- **Simplify the process of validating and provisioning resources for business units.** Datacenter administrators can use the self-service portal to assign resources based on business unit requests.
- **Control the change process for these resources.** Changes to resources follow a request-and-approve life cycle, and the requests remain on record in the database.

Assumptions

System Center Virtual Machine Manager 2008 R2 makes a number of features and functionality possible. However, this document is scoped to only include using System Center Virtual Machine Manager 2008 R2 with stand-alone Hyper-V hosts as a basis for managing the automated provisioning of the virtual machines on these hosts with the Self-Service Portal v2.0. The document further includes server consolidation using physical-to-virtual and virtual-to-virtual methods.

Microsoft System Center Virtual Machine Manager is designed to take advantage of the latest features and benefits found in the Windows® Server and Microsoft® System Center Family. With this in mind System Center Virtual Machine Manager will only install on Windows Server® 2008 or Windows Server® 2008 R2 and will be installed using Microsoft® SQL Server® 2008 to accommodate the SSP 2.0 requirements.

MICROSOFT WINDOWS SERVER ® 2008 R2 WITH HYPER-V™

Windows Server 2008 R2 and Hyper-V System Requirements

This section outlines the supported operating systems and minimum system requirements for a supported Windows Server® 2008 R2 server running the Hyper-V™ role. Subsequent sections of the document provide detailed installation and configuration procedures.

Supported Host Operating Systems:

- Windows Server® 2008 R2 Standard Edition x64 with Hyper-V™
- Windows Server® 2008 R2 Enterprise Edition x64 with Hyper-V™
- Windows Server® 2008 R2 Datacenter Edition x64 with Hyper-V™

> **Note**
>
> The Standard Edition does not support Hyper-V™ High Availability configurations.
>
> This document does not address Microsoft® Hyper-V™ Server R2, which now supports high-availability configurations.

Intel Processor Requirements:

- x64 Processor Architecture
- Support for Hardware Execute Disable
- Intel® VT Hardware Virtualization

AMD Processor Requirements

- x64 Processor Architecture
- Support for Hardware Execute Disable
- AMD-V® Hardware Virtualization

Minimum CPU Speed: 1.4 GHz
RAM: Minimum of 512 MB of RAM
Required Available Disk Space: 10 GB of available hard disk space

> **Note**
>
> Computers with more than 16 GB of RAM will require more disk space for paging and dump files.

Hyper-V R2 Host Limitations

Functionality	Windows Server® 2008 R2 Standard Edition	Windows Server® 2008 R2 Enterprise Edition	Windows Server® 2008 R2 Datacenter Edition
Logical Processor Support	64 LP	64 LP	64 LP
Physical Memory Support	Up to 32 GB	Up to 1 TB	Up to 1 TB
Max # of VMs	8 V-Procs per LP or 384 VMs, whichever is lower	8 V-Procs per LP or 384 VMs, whichever is lower	8 V-Procs per LP or 384 VMs, whichever is lower
VM Licensing	1 Free Per License	4 Free Per License	Unlimited

> **Note**
>
> These limitations are for the Hyper-V R2 role only, not the Windows Server Operating System.

Hyper-V R2 Guest Limitations

- x86 or x64 operating systems
- Up to 4 logical processors
- Up to 64 GB of RAM per guest
- Up to 4 IDE devices
- Up to 4 SCSI controllers supporting up to 64 disks each
- Up to 4 legacy network adapters
- Up to 8 synthetic network adapters

Supported Operating System	Virtual

	Processors		
	1	2	4
Windows Server® 2008 R2	x	x	x
Windows Server® 2003 x86x64 w/ SP2	x	x	
Windows® 2000 Server & Advanced Server w/ SP4	x		
Windows® HPC Server 2008	x	x	x
SUSE® Linux Enterprise Server 10 x86x64 w/ SP1/SP2	x		
Red Hat® Enterprise Linux	x	x	x
Windows 7	x	x	x
Windows Vista® x86/x64 w/ SP1	x	x	
Windows® XP Pro x64 w/ SP2 & x86 w/ SP3	x	x	
Windows® XP Pro x86 w/ SP2	x		

Standalone Host Architecture Patterns

Hyper-V Single Host Server Architecture

The single host server architecture pattern is illustrated below. The architecture consists of a single host server running Windows Server 2008 R2 with Hyper-V that runs a number of virtual machine guests. This pattern provides server consolidation but does not provide high availability. The host server is a single point of failure and this architecture will necessitate a Save State or Power Off of the virtual machine guests should the host server require maintenance or reboot.

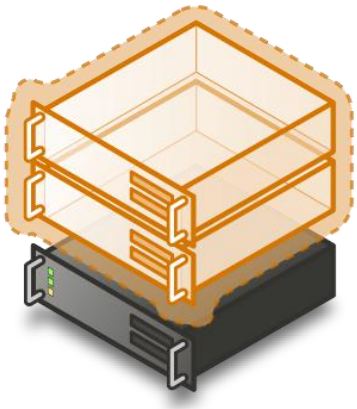

Windows Server 2008 with Hyper-V Host Server

Standalone Host Storage Architecture

The storage design that is utilized for the host server architecture has a major impact on host and guest performance. Storage performance is a complicated mix of drive, interface, controller, cache, protocol, SAN, HBA, driver, and operating system considerations. The overall performance of the storage architecture is typically measured in terms of Maximum Throughput, Maximum IO operations per second (IOPS), and Latency or Response Time. While each of the three factors is important, IOPS and Latency are the most relevant to server virtualization.

This section outlines the various storage architecture considerations and provides recommendations for each of the overall architecture patterns.

Storage Connectivity

Individual disks and storage arrays can be accessed by host servers in three different ways: Direct Attached Storage, iSCSI Storage Area Networks, and Fibre Channel Storage Area Networks.

Direct Attached Storage

Direct attached storage typically consists of hard drives internal to the host server itself or hard drives housed in a dedicated storage array attached directly to the server via SCSI, SAS, or eSATA connections. The host server utilizes an internal SCSI, SAS, or SATA controller card to enable the server to access the storage and enable various RAID levels. The storage array is typically dedicated to the individual server.

iSCSI Storage Area Network

iSCSI is an increasingly popular storage networking architecture that enables use of the SCSI protocol over TCP/IP network infrastructures. iSCSI enables the use of standard Ethernet networking components such as NICs, Switches, and Routers to build a storage area network. Typically iSCSI SANs are less expensive to implement than traditional Fibre Channel SANs. The storage array used in iSCSI architecture is usually a low- to mid-range array that is shared by multiple host servers. Redundant, dedicated gigabit Ethernet NICs are recommended for iSCSI connectivity.

Fibre Channel Storage Area Network

Fibre Channel storage area networks provide high speed, low latency connectivity to storage arrays. Host Bus Adapters (HBAs) are utilized by the host servers to connect to the Fibre Channel SAN via switches and directors. Fibre Channel SANs are typically used in concert with mid to high end storage arrays, which provide a multitude of features such as RAID, disk snapshots, multipath IO, and so on.

> **Recommendation**

> Direct attached storage or iSCSI SAN storage is recommended for the Single Host Server architecture pattern.
>
> For performance and security reasons, it is strongly recommended that iSCSI SANs utilize dedicated NICs and switches that are separate from the LAN.

Drive Types

The type of hard drive utilized in the host server or the storage array the host servers will have the most significant impact on the overall storage architecture performance. The critical performance factors for hard disks are the interface architecture (for example, U320 SCSI, SAS, SATA), the rotational speed of the drive (7200, 10k, 15k RPM), and the average latency in milliseconds. Additional factors, such as the cache on the drive, and support for advanced features, such as Native Command Queuing (NCQ), can improve performance.

As with the storage connectivity, high IOPS and low latency are more critical than maximum sustained throughput when it comes to host server sizing and guest performance. When selecting drives, this translates into selecting those with the highest rotational speed and lowest latency possible. Utilizing 15k RPM drives over 10k RPM drives can result in up to 35% more IOPS per drive.

Use the below information to evaluate the cost/performance tradeoffs.

SCSI

SCSI drives are rapidly being replaced by SATA, SAS, and Fibre Channel drives. SCSI drives are not recommended for new host server architectures; however, existing servers with U320 SCSI drives can provide excellent performance characteristics.

SATA

SATA drives are a low cost and relatively high performance option for storage. SATA drives are available primarily in the 1.5 GB/s and 3.0 GB/s standards (SATA I and SATA II) with a rotational speed of 7200 RPM and average latency of around 4 ms. There are a few SATA I drives that operate at 10k RPM and average latency of 2 ms that can provide an excellent low cost storage solution.

SAS

SAS drives are typically much more expensive than SATA drives but can provide significantly higher performance in both throughput, and more importantly, low latency. SAS drives typically have a rotational speed of 10k or 15k RPM with an average latency of 2 to 3 ms.

Fibre Channel

Fibre Channel drives are usually the most expensive and typically have similar

performance characteristics to SAS drives but use a different interface. The choice of Fibre Channel or SAS drives is usually determined by the choice of storage array. As with SAS, they are typically offered in 10k and 15k RPM variants with similar average latencies.

If you are using a Fibre Channel SAN, ensure that the switch and director infrastructure is sized to handle the large amount of storage I/O that will be generated from the consolidated servers.

> **Recommendation**
>
> SATA drives with a minimum of 7200 RPM are recommended for the Single Host Server architecture pattern, though SAS 10k or 15k RPM drives are preferred.

Disk Redundancy Architecture

Redundant Array of Inexpensive Disk (RAID) is strongly recommended for all Hyper-V host storage. By definition, Hyper-V hosts run and store data from multiple workloads. RAID is necessary to ensure that availability is maintained during disk failure. In addition, if properly selected and configured, RAID arrays can provide improvements in overall performance.

RAID 1

RAID 1 is disk mirroring. Two drives store identical information so that one is a mirror of the other. For every disk operation, the system must write the same information to both disks. Because dual write operations can degrade system performance, many employ duplexing, where each mirror drive has its own host adapter. While the mirror approach provides good fault tolerance, it is relatively expensive to implement because only half of the available disk space can be used for storage, while the other half is used for mirroring.

RAID 5

Also known as striping with parity, this level is a popular strategy for low- or mid-range storage systems. RAID 5 stripes the data in large blocks across the disks in an array. RAID 5 writes parity data across all the disks in the RAID 5 set. Data redundancy is provided by the parity information. The data and parity information is arranged on the disk array so that the two types of information are always on different disks. Due to the nature of its parity algorithm, each write request incurs 3 actual writes to the disks, decreasing write performance. Striping with parity can offer better performance than disk mirroring (RAID 1). However, when a stripe member is missing, read performance is decreased (for example, when a disk fails). RAID 5 is a less expensive option because it utilizes drive space more efficiently than RAID 1.

RAID 10 (RAID 1+0)

This level is also known as mirroring with striping. RAID 10 uses a striped array of disks that are then mirrored to another identical set of striped disks. For example, a striped array can be created by using five disks. The striped array of disks is then mirrored using another set of five striped disks. RAID 10 provides the performance benefits of disk striping with the disk redundancy of mirroring. RAID 10 provides the highest read-and-write performance of any one of the other RAID levels, but at the expense of using twice as many disks.

RAID 50 (RAID 5+0)

This is a nested RAID level combining the block-level stripping of RAID 0 with the parity of RAID 5. It can be thought of as a RAID 0 array consisting of multiple RAID 5 arrays. This level improves upon the write performance of RAID 5 and provides better fault tolerance than a single RAID level. The specific configuration and number of disks will determine the actual availability and performance characteristic of this RAID level. This RAID type is becoming a common feature on even low-end storage devices.

Other RAID levels may offer additional fault tolerance or performance enhancements. These levels generally are proprietary systems. For more information about these types of RAID systems, contact your storage hardware vendor.

> **Recommendation**
>
> RAID 1 is recommended for the system volume in all host server architecture patterns.
>
> RAID 1 or RAID 10 is recommended for the data volumes in the Single Host Server architecture pattern.
>
> RAID 5 and 50 are generally not recommended for virtualization environments due to their inherent write performance penalties.

Storage Controller Architecture

The storage controller is either a server add-in card, such as a SCSI or SAS controller, or a component of a mid- to high-end storage array. The storage controller provides the interface between the disk drives and the server or storage area network. The design factors that impact storage controller performance include the interface or HBA type, the amount of cache, and the number of independent channels.

Disk Controller or HBA Interface

The disk controller interface determines the types of drives that can be utilized as well as the speed and latency of the storage I/O. The table below summarizes the

most commonly utilized disk controller interfaces.

Architecture	Throughput (theoretical max Megabyte/sec)
iSCSI (Gigabit Ethernet)	125 MB/s
Fibre Channel (2 GFC)	212.5 MB/s
SATA (SATA II)	300 MB/s
SCSI (U320)	320 MB/s
SAS	375 MB/s
Fibre Channel (4 GFC)	425 MB/s
Fibre Channel (8 GFC)	850 MB/s
iSCSI (10 Gigabit Ethernet)	1250 MB/s

> **Recommendation**
>
> SATA II or SAS are recommended for the Single Host Server architecture pattern, with SAS being the preferred option.

Controller Cache

Storage controller cache can improve performance during burst transfers or when the same data is accessed frequently by storing it in the cache memory, which is typically several orders of magnitude faster than the physical disk I/O.

> **Recommendation**
>
> When comparing similar storage controllers or selecting storage controller options, select those that have a larger and faster cache memory**.**

Controller Channels

The number of internal and external channels that a storage controller has can substantially impact overall storage performance. Multiple channels increase the amount of simultaneous read and write IO operations (IOPS) that can be performed, especially when using advanced RAID arrays.

> **Recommendation**
>
> Utilize a minimum of a 2-channel storage controller for the Single Host Server architecture pattern. Utilize one channel for the RAID 1 system partition and

the other for the RAID 10 data partition(s).

Utilize the best practices from your storage vendor to distribute the RAID 10 mirrors and stripes across the multiple channels for maximum performance.

Note

This section addressed storage system considerations and recommendations. The next section addresses server considerations such as processors, RAM, and I/O.

HOST SERVER ARCHITECTURE

The host server architecture is a critical component of the virtualized infrastructure, as well as a key variable in the consolidation ratio and cost analysis. The ability of the host server to handle the workload of a large number of consolidation candidates increases the consolidation ratio and helps provide the desired cost benefit.

The "sweet spot" is typically in the two- to four-socket servers running the highest or second highest CPU speed with multi-core processors.

Note

There are programs available for assisting customers in selecting hardware, but they do not address sizing or configuration of that hardware: Windows Server Catalog contains all servers, storage, and other hardware devices that are certified for Windows Server 2008 R2 and Hyper-V.

Windows Server Catalog:

Go to www.windowsservercatalog.com.

- Click Certified Servers.
- Then click Hyper-V (bottom-left).

System Architecture

The system architecture of the host server refers to the general category of the server hardware itself. Examples include rack mounted servers, blade servers, and large symmetric multiprocessor servers (SMP). The primary tenet to consider when selecting system architectures is that each Virtual Server host will contain multiple guests with multiple workloads. Processor, RAM, Storage, and Network capacity are critical, as well as high I/O and low latency. It is critical to ensure that the host server is able to provide the required capacity in each of these

categories.

Standard Rack Mounted Servers

The most common system architecture is a standard rack mounted server. Typically found in 2U or 4U models, these servers typically contain 2 to 4 CPU sockets, 2 to 8 PCI-E or PCI-X slots, and 4 to 6 hard disk bays. Rack mounted servers are excellent choices for Hyper-V hosts due the low cost of commodity 2- and 4-socket servers and their inherent scalability and expandability through additional NIC and HBA slots.

> **Recommendation**
>
> Rack mounted Intel- or AMD-based servers are recommended for any of the host server architecture patterns.

Blade Servers

Blade servers have dramatically increased in popularity and capability due to the ever increasing need for capacity and server density. Blade server architectures are a primary area of R&D for the server manufacturers, resulting in a significant amount of innovation in this space. The downside of blade servers is limited standards and interoperability between manufacturers and, in some cases, within the same manufacturer when they change their blade chassis architecture.

The processor and space density benefits of blade servers initially came at the expense of expandability and the quantity of NICs and HBAs that can be supported in the first several generations of blade servers.

Recently, the advent of blade architectures where each blade contains 8 to 16 cores, up to 64 GB of RAM, and most importantly, 6 or more IO interfaces has eliminated many disadvantages that previously mitigated against using blade server architectures for virtualization.

The network and storage I/O that is required to support the desired number of guests on each host server must be considered carefully to ensure that each host server is running on a blade and the blade chassis itself can support it.

The host server architecture must be considered when evaluating blade server system architectures. If an iSCSI storage system is planned, two additional dedicated NICs are required for access to the storage and redundancy. Finally, at least two NICs should be dedicated to network I/O. The number of NICs required per host can easily expand from 4 to 6 or more NICs. This is frequently beyond the number supported by many blade servers.

> **Warning**

> Microsoft does not support the use of NIC teaming software. Support for these third-party technologies must be provided by the vendor.

> **Recommendation**
>
> Blade servers are recommended for any of the host server architecture patterns. Careful analysis of network and I/O requirements must be performed to ensure that each blade server and the blade chassis itself can handle the total I/O required.
>
> Blade servers are also excellent candidates if multiple departments or business groups are going to have dedicated Hyper-V hosts or small groups of hosts.

Large SMP Servers

For the purposes of this document, large SMP servers are defined as those that have 8 or more CPU sockets. At the very high end, Windows Server 2008 R2 Datacenter Edition on 64-bit hardware can support servers with up to 64 CPU sockets and 2 TB of RAM. Many of these very high end servers include advanced features such as hardware partitioning, hot-add of resources, hot spare components, and so on. Hardware of this capacity has the potential to host hundreds of virtual machine guests.

While providing an excellent consolidation ratio, large SMP servers are typically much more expensive than the commodity 2- and 4- socket servers described earlier. A fully populated 32-socket large SMP machine can easily cost over $500,000 USD, compared to a 4-socket commodity server costing $30,000 USD.

A large SMP server or large SMP server cluster may be appropriate if a very large number of servers will be consolidated, if the organization has operational experience with large "mainframe" class servers, or if they have already standardized on large SMP hardware.

> **Recommendation**
>
> Large SMP servers are only recommended for organizations that have excellent operational experience with large mission-critical servers or those that have already standardized on this platform.

Processor Architecture

Windows Server 2008 R2 with Hyper-V requires x64 processor architecture from Intel or AMD, as well as support for hardware execute disable and hardware virtualization such as Intel VT or AMD-V.

Both Intel and AMD provide a wide range of processors that are appropriate for host servers. The industry competition between the two is very tight and at any one time; one may have a performance advantage over the other. Regardless of which manufacturer is chosen, several performance characteristics are important.

The number of processor cores is a key performance characteristic. Windows Server 2008 R2 with Hyper-V makes excellent use of multi-core processors, so the more cores the better. Another important characteristic is the processor clock speed, which is the speed at which all cores in the processor will operate. It's important because it will be the clock speed of all of the guest virtual machines. This is a key variable in the consolidation ratio because it impacts the amount of candidates that the host server can handle AND the speed at which those guests will operate. As an example, choosing 2 GHz processor rather than a 3 GHz processor on a server that will host 20 guests means that all of those guests will run only at 2 GHz.

At a lower level of detail, the server processor architectures make design choices in terms of the type and quantity of processor cache, memory controller architecture, and bus/transport architecture. A detailed analysis of these factors is beyond the scope of this document.

> **Recommendation**
>
> x64 processor architectures are required for all Hyper-V host server architecture. If you are purchasing new servers, we recommend working with your server vendor to ensure that the selected hardware is capable of running Windows Server 2008 R2 and Hyper-V, and that it is validated for Windows Server 2008 R2 failover clustering. For new servers, we recommend selecting the maximum number of cores per processor available and choosing the fastest or second fastest clock speed available.

Memory Architecture

Once the system architecture and processor architecture choices are made, there are relatively few options remaining for memory architecture because it is usually predetermined by the manufacturer/system/processor combination. The memory architecture choices that remain are typically quantity, speed, and latency. For Hyper-V, the most important memory architecture choice is the quantity of RAM. Most consolidated workloads (that is, individual guest virtual machines) will require at least 512 MB to 1 GB of RAM or more. Since most commodity four-socket servers can only cost effectively support between 32 and 128 GB of RAM, this is frequently the limiting factor in host server capacity.

The quantity of RAM is a more important factor than RAM speed or latency.

Once the maximum amount of RAM that is cost effective is determined, if there is a remaining choice between speed and latency, choosing the memory with lower latency is recommended.

> **Recommendation**
>
> Given the system and processor architectures already selected, we recommend utilizing the maximum amount of RAM that can be cost effectively added to the host system. Typically, there is a price point where the cost of moving to the next DIMM size (that is, 2 GB DIMMs to 4 GB DIMMs) is more than twice the cost, and in some cases, it approaches the cost of an entire server. We recommend fully populating the server up to that price point. For example, if the server has 8 DIMM slots and 4 GB DIMMs are much more than twice the cost of 2 GB DIMMs, we recommend fully populating the server with 2 GB DIMMs and considering a second host server if additional capacity is required.
>
> For all host server architecture, a minimum of 16 GB of RAM is recommended

Network Architecture

The network architecture of the host server is a frequently overlooked topic in host server sizing because Gigabit Ethernet NICs are now very inexpensive and most servers have at least two built in. The topic is important, however, because it is directly impacted by the host server architecture pattern selected. As mentioned previously, if an iSCSI storage architecture is being utilized, NICs will need to be dedicated to storage I/O traffic. Gigabit Ethernet is a high-speed network transport, though a host server with a large number of guests may require greater than Gigabit speed, thus requiring additional NICs. Finally, it is recommended that each host server have a NIC dedicated to the host itself for network I/O and management.

As described earlier, a fairly large number of NICs per host server may be required. This is the one factor that can mitigate against blade servers in some instances. Recently, 10-Gigabit Ethernet has become commonly available and is starting to drift lower in price, similar to the way Gigabit Ethernet has done over the years. The ability for servers to utilize 10-Gigabit Ethernet NICs is a significant factor in increasing the consolidation ratio.

> **Recommendation**
>
> Use multiple NICs and multi-port NICs on each host server.
>
> One NIC dedicated to the host machine only for management purposes

> One or more NICs dedicated to the guest virtual machines (use 10 gpbs NICS for highest consolidation)
>
> Two or more NICs dedicated to iSCSI with MPIO
>
> Dedicate at least one NIC/Port on each host server for guest virtual machine network I/O. For maximum consolidation ratio, utilize one or more 10-Gigabit Ethernet NICs to virtual machine network I/O.

Host Bus Adapter Architecture

The disk storage for all guest virtual machines is one or more VHD files housed on the storage system being utilized by the host server. Host storage I/O, in addition to the system, processor, memory, and network architectures described earlier, is the final major component of host server sizing. Hyper-V I/O consists of a large number of read and write IOPS to the storage system due to the large number of guests running on each server and their various workloads.

If direct attached storage is being utilized, a SATA II or SAS RAID controller internal to the server is recommended as described earlier. If a storage array and SAN are being utilized, host bus adapters (HBAs) are required in the host server. The HBA provides access to the storage array from the host server. The storage connectivity is a critical component for high availability and performance.

Recommendation

> Utilize at least one Fibre Channel HBA or one dedicated NIC for iSCSI with the Single Host Server architecture pattern.
>
> If using Fibre Channel, utilize 4 GFC or 8 GFC HBAs.
>
> If using iSCSI, utilize MPIO in a load-balancing configuration for maximum throughput.

Operating System Architecture

Operating System Version

The choice of operating system for the Hyper-V hosts is important from a support and performance perspective, as well as an overall cost perspective. As mentioned earlier, an x64 version of Windows Server is required in all scenarios.

Also, consider virtualization use rights when choosing the operating system version. Certain versions of Windows Server 2008 R2 (namely Standard, Enterprise, and Datacenter editions) include "virtualization use rights," which is the right and license to run a specified number of Windows-based virtual machines. Windows Server® 2008 R2 Standard edition includes use rights for

one running virtual machine. Windows Server® 2008 R2 Enterprise Edition includes use rights for up to four virtual machines. This does not limit the number of guests that the host can run; it means that licenses for four Windows guests are included. To run more than four you simply need to ensure you have valid Windows Server licenses for the additional virtual machines.

Windows Server® 2008 R2 Datacenter Edition includes unlimited virtualization use rights, which allows you to run as many guests as you like on the physical server running Windows Server 2008 R2 Datacenter edition.

> **Recommendation**
>
> Use Windows Server® 2008 R2 Enterprise or Windows Server® 2008 R2 Datacenter editions for all Hyper-V hosts. Work with your account executive to determine the point at which Datacenter server becomes cost effective once you have completed the host server sizing methodology and have determined how many guests you will run on each host.
>
> Reference Microsoft Licensing for Virtualization.

Hyper-V Architecture

Virtual Machine Guests

Hyper-V greatly increases the scalability of guest virtual machines in comparison to Virtual Server 2005. Hyper-V guests, when utilizing a supported operating system, are able to support the following options:

> **Note**
>
> Verify the support level for multiple processors, large RAM, and so on, for each operating system planned to utilize those settings.

Since Hyper-V supports such large guest virtual machines, there is a much larger set of viable workloads that can be consolidated, including multi-processor, multi-core servers, servers with large disk or IO requirements, and so on.

While Hyper-V can support large guests, in general it is prudent to only configure each guest with the resources needed. This ensures that resources are available for other guests or future expansion. For example, it is not recommended to make all guests utilize four logical processors if that is not specifically needed for the guests. Additional resources such as processors, RAM, and so on, can easily be added if needed.

> **Recommendation**

> Configure guests to utilize only the resources needed to achieve the desired performance and maximum consolidation ratio.

The diagram below illustrates a guest virtual machine configured with a moderate amount of resources such as 4 logical processors, 4 GB of RAM, multiple SCSI controllers, and multiple network adapters running Windows Server 2008. In this example, the guest includes an IDE boot disk (VHD) and four SCSI pass-through disks. Guest storage architecture is detailed in the next section.

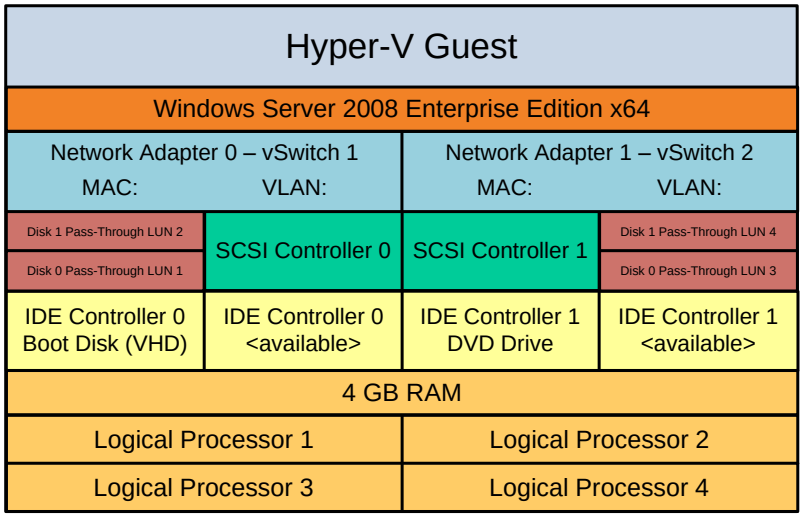

Virtual Machine Storage

Volumes and Partitions

Guests running on Windows Server 2008 R2 and Hyper-V host servers benefit from many of the same disk I/O performance tuning techniques as servers running Microsoft® SQL Server® or Microsoft® Exchange Server. Dedicating a high speed LUN to the operating system and placing virtual hard disk files (.VHDs) and virtual machine configuration files on separate high speed LUNs are recommended. Segregating disk I/O onto separate physical spindles may still be appropriate depending on the workload's performance characteristics. Please reference the application-specific guidance for disk I/O segregation recommendations.

Hyper-V also provides the option to use pass-through disks that provide the guest with direct access to a LUN without the disk being presented to the host. This feature may be a good fit when considering reallocating storage. For example, when VM data reaches a certain size it makes more sense to re-map the LUN rather than copy the data. At that point, consider pass-through disks.

If you are using a storage array, confirm with your storage vendor the appropriate track and sector values for your storage system and use the Diskpart.exe tool to verify that your disk tracks are sector-aligned. In most cases with Windows Server 2008 R2, this is not necessary, but it should be verified with

your storage vendor.

> **Recommendation**
>
> Use separate physical disks and LUNs for VM Guest operating system VHD files and data volumes.
>
> Segregate disk I/O as per the Guest application's disk tuning guidance.
>
> Utilize NTFS for all host server volumes.
>
> In operating systems prior to Windows Server 2008, align the Guest disk sectors as per http://support.microsoft.com/kb/929491

Periodically defragmenting, pre-compacting, and compacting the VHD files on the guest and defragmenting the volumes on the host will help ensure optimal disk I/O performance. If using only fixed-size VHDs, host-level defragmentation of the volumes hosting VHDs is not necessary because the disk sectors are pre-allocated in a contiguous fashion at the time of creation.

> **Information**
>
> Defragmentation of the host can be performed by using the built-in Microsoft Windows® disk defragmentation utility. Detailed instructions for VHD defragmentation, pre-compaction, and compaction can be found in this article:
> http://vscommunity.com/blogs/virtualzone/archive/2007/01/17/three-steps-to-vhd-compaction-with-virtual-server-2005-r2-sp1.aspx

Virtual Hard Disks (VHD)

Virtual hard disks encapsulate a guest's hard disk inside of a VHD file, which is placed on storage that is accessible to the host server. Utilizing virtual hard disks provides benefits such as the ability to dynamically expand the disk, the ability to take snapshots of the disk, portability in terms of moving the disk to a different server, and so on. There are three forms of virtual hard disks.

Dynamically Expanding Disks

Dynamically expanding virtual hard disks provide storage capacity as needed to store data. The size of the .vhd file is small when the disk is created and grows as data is added to the disk. The size of the .vhd file does not shrink automatically when data is deleted from the virtual hard disk. However, you can compact the disk to decrease the file size after data is deleted by using the Edit Virtual Hard Disk Wizard.

Fixed Size Disks

Fixed virtual hard disks provide storage capacity by using a .vhd file that is in the size specified for the virtual hard disk when the disk is created. The size of the .vhd file remains 'fixed' regardless of the amount of data stored. However, you can use the Edit Virtual Hard Disk Wizard to increase the size of the virtual hard disk, which increases the size of the .vhd file. By allocating the full capacity at the time of creation, fragmentation at the host level is not an issue (fragmentation inside the VHD itself must be managed within the guest).

Differencing Disks

Differencing virtual hard disks provide storage to enable you to make changes to a parent virtual hard disk without altering that disk. The size of the .vhd file for a differencing disk grows as changes are stored to the disk.

> **Recommendation**
>
> For production environments, utilize Fixed Size disks, which provide better performance and ease the monitoring of storage availability. Utilizing fixed disks allocates the full size of the disk upon creation.
>
> In Hyper-V R2, the performance of Dynamically Expanding disks (including Snapshots / .AVHDs, and Differencing Disks) has increased dramatically and are now viable options for production use. However, they carry other risks such as storage oversubscription and fragmentation, so use with caution.

Pass-Through Disks

Hyper-V enables virtual machine guests to directly access local disks or SAN LUNs that are attached to the physical server without requiring the volume to be presented to the host server. The virtual machine guest accesses the disk directly (utilizing the disk's GUID) without having to utilize the host's file system. Given that the performance difference between Fixed-Disk and Pass-through Disks is now negligible, the decision is now based on manageability. For instance, if the data on the volume will be very large (hundreds of gigabytes), a VHD is hardly portable at that size given the extreme amounts of time it takes to copy. Also, bear in mind the backup scheme. With pass-through disks, the data can only be backed up from within the Guest.

When utilizing pass-through disks, there is no VHD file created; the LUN is used directly by the guest. Since there is no VHD file, there is no dynamic sizing capability or snapshot capability.

> **Recommendation**
>
> Use pass-through disks only in cases where absolute maximum performance is required and the loss of features such as snapshots and

> portability is acceptable. Since the performance difference between pass-through and fixed-disks is minimal there should be very few scenarios where pass-through disks are required.

Disk Access Options

Virtual machine guests can access storage by utilizing three mechanisms: IDE, SCSI, and iSCSI. When configuring IDE or SCSI disks for a guest, you can choose either a VHD or pass-through disk configuration utilizing any storage connected to the host server (that is, disks direct attached to the host, SAN LUNs presented to the host, or iSCSI LUNs presented to the host).

While diagrammed separately, the various options can be combined and used together.

In each of the diagrams, the blue disks represent storage mounted by the host, which holds VHD files used by the guests. The orange disks represent storage accessed directly by the guests using either pass-through disks (using either IDE or SCSI virtual controllers) or directly connecting to iSCSI LUNs that is presented to the guest.

In this diagram, direct attached storage such as SAS, SCSI, or SATA disks is utilized:

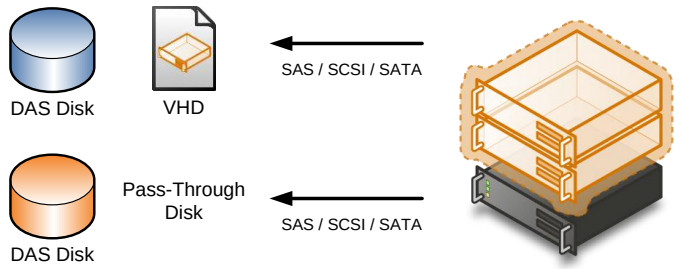

In this diagram, Fibre Channel SAN-based storage is utilized:

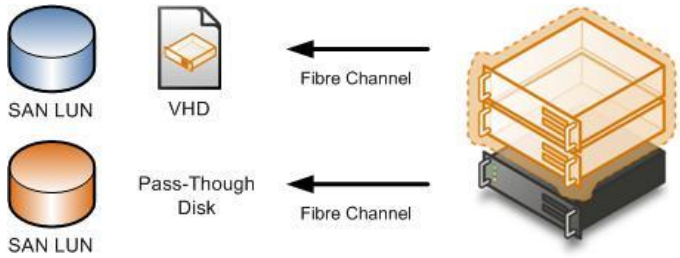

> **Note**
>
> Hyper-V guests can only boot from IDE disks. The Hyper-V guest BIOS

supports two IDE controllers, each supporting up to two disks for a maximum of four IDE disks per guest.

Hyper-V guests support up to four SCSI controllers, each supporting up to 64 disks for a total of up to 256 SCSI disks per guest.

Unlike Virtual Server 2005 R2, once the Hyper-V integration components are installed in a guest, there is no performance difference between IDE or SCSI connected virtual disks.

Recommendation

Utilize an IDE connected disk as the boot disk. Add an SCSI controller and SCSI connected disks for guest data volumes.

In Hyper-V R2, Guest Disks can be hot-added via the virtual SCSI controller. Therefore, consider pre-creating a SCSI controller device on all VMs for hot-add VHD capability.

Hyper-V can also utilize iSCSI storage by directly connecting to iSCSI LUNs utilizing the guest's virtual network interface cards (NICs). Guests cannot boot from iSCSI LUNs accessed through the virtual NICs without utilizing a third-party iSCSI initiator.

In this diagram, iSCSI storage is utilized. With iSCSI, a third access scenario is added, which is direct iSCSI access utilizing the network connectivity of the guest.

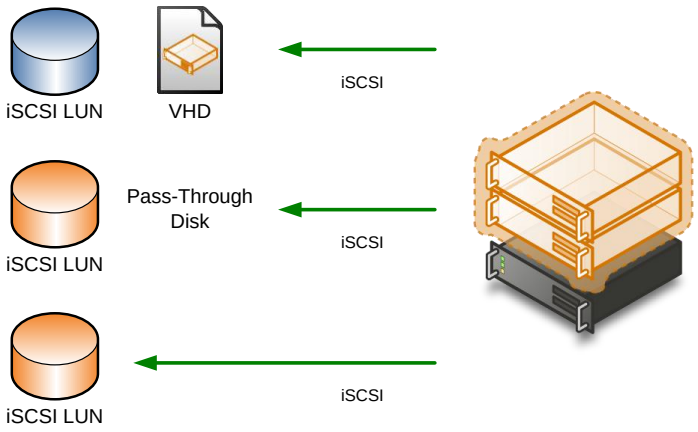

Note

Do not confuse iSCSI LUNs presented to the host and then utilized by the guest with iSCSI LUNs presented directly to the guest. In the former, access to the iSCSI LUN is provided via the host's network connectivity. In the latter, access to the iSCSI LUN is provided via the guest's network

connectivity. The next section describes these options.

Recommendation

If you are using iSCSI, ensure that a separate physical and virtual network is utilized for access to the iSCSI storage to obtain acceptable performance, including cabling and switching.

If you are utilizing iSCSI LUNs presented to the host, this means having dedicated physical NIC(s) for connectivity to the iSCSI storage.

Consider using Jumbo Frames on both the Guest's and Host's storage providing NICs for performance improvements.

If you are utilizing iSCSI LUNs directly presented to the guests, this means having dedicated Physical NIC(s) connected to the host, dedicated virtual switch (es) bound to the iSCSI physical NIC(s), and dedicated virtual NIC(s) in the guests bound to the iSCSI virtual switch. The end result would be a guest with two or more virtual NICs configured one for LAN connectivity and one or more for iSCSI connectivity.

Virtual Networks

You can create many virtual networks on the server running Hyper-V to provide a variety of communications channels. For example, you can create networks to provide the following:

- Communications between virtual machines only. This type of virtual network is called a private network.
- Communications between the Host server and virtual machines. This type of virtual network is called an internal network.
- Communications between a virtual machine and a physical network by creating an association to a physical network adapter on the host server. This type of virtual network is called an external network.

You can use Virtual Network Manager to add, remove, and modify the virtual networks. Virtual Network Manager is available from Hyper-V Manager MMC. The network types are illustrated below.

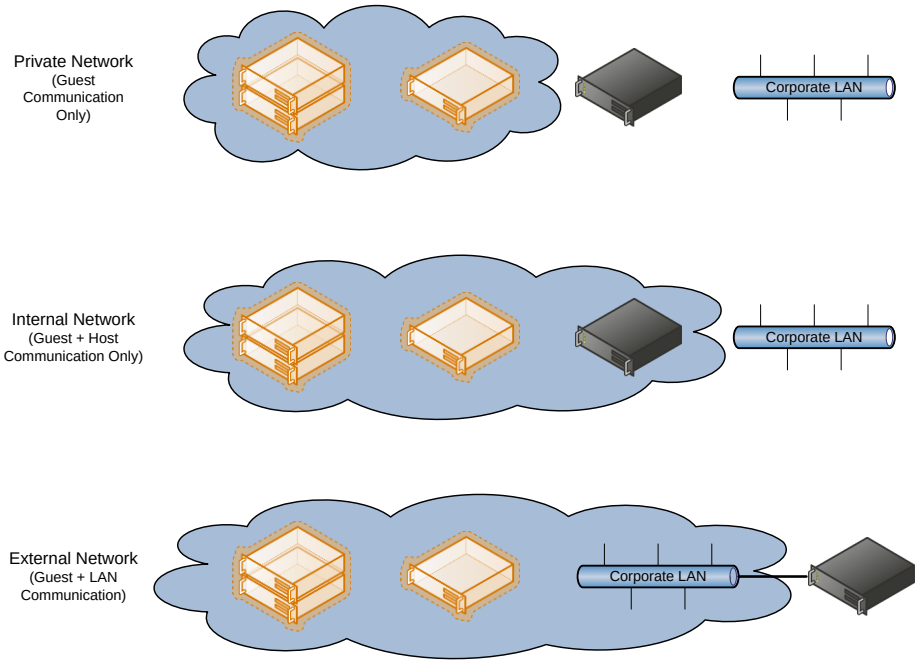

When creating an External network in Hyper-V, a virtual network switch is created and bound to the selected physical adapter. A new virtual network adapter is created in the parent partition and connected to the virtual network switch. Child partitions can be bound to the virtual network switch by using virtual network adapters. The diagram below illustrates the architecture.

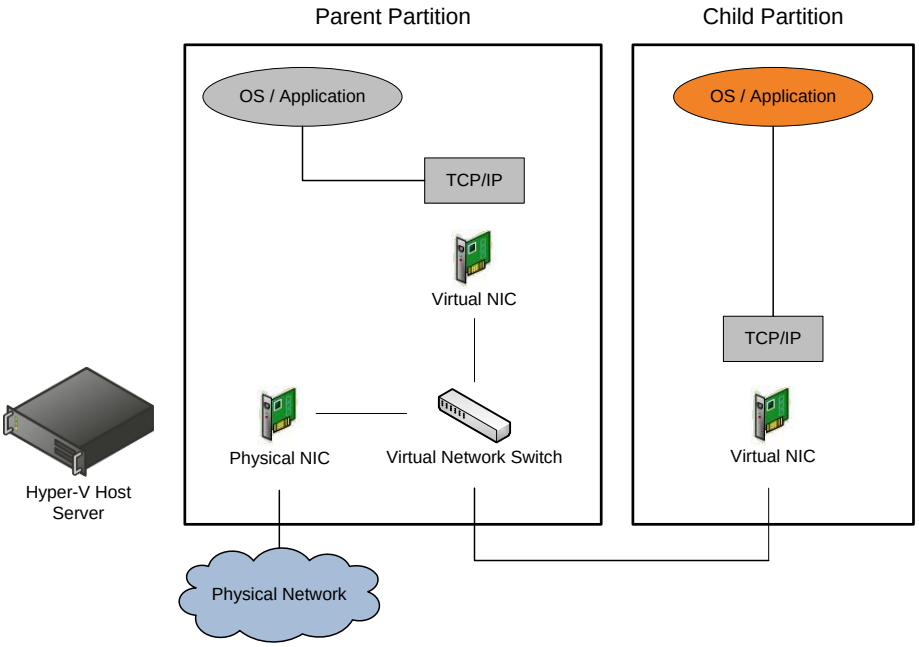

In addition to the above scenarios, Hyper-V also supports the use of VLANs and VLAN IDs with the virtual network switch and virtual network adapters. Hyper-V leverages 802.1q VLAN *trunking* to achieve this objective. To utilize this functionality, a virtual network switch must be created on the host and bound to a physical network adapter that supports 802.1q VLAN tagging. VLAN IDs are configured in two places:

- The virtual network switch itself, which sets the VLAN ID the parent partition's virtual network adapter will use
- The virtual network adapter of each guest, which sets the VLAN ID the guest will use

The diagram below illustrates an example of using a single physical NIC in the host that is connected to an 802.1q trunk on the physical network carrying three VLANs (5, 10, 20). The design objectives in this example are:

- An 802.1q trunk carrying 3 VLANs (5, 10, 20) is connected to a physical adapter in the host.
- A single virtual switch is created and bound to the physical adapter.
- The VLAN ID of the virtual switch is configured to 5, which would allow the virtual NIC in the parent to communicate on VLAN 5.
- The VLAN ID of the virtual NIC in Child Partition #1 is set to 10, allowing it to communicate on VLAN 10.
- The VLAN ID of the virtual NIC in Child Partition #2 is set to 20, allowing it to communicate on VLAN 20.

The expected behavior is that there is a single virtual switch; the parent and two children can only talk on their respective VLANs, and they can't talk to each other.

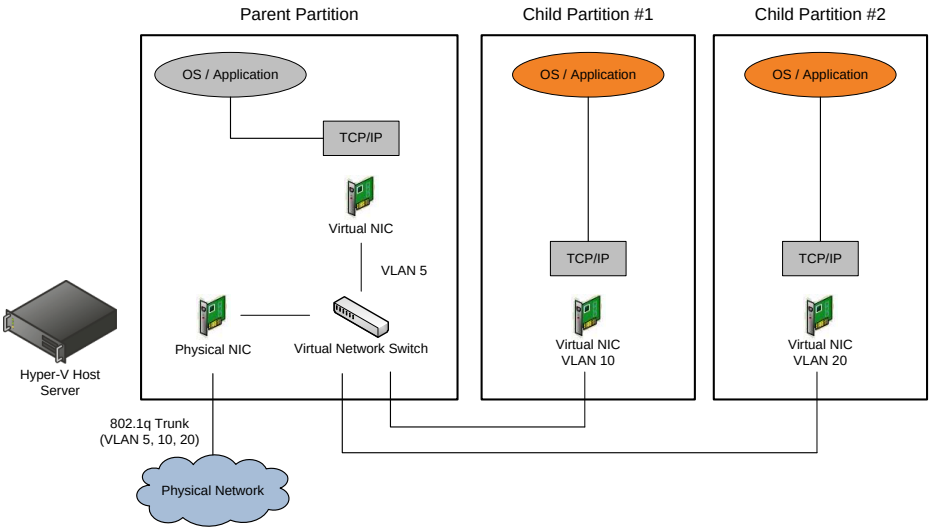

Security Considerations

Microsoft Hyper-V was designed to minimize the attack surface on the virtual environment. The Hypervisor itself is isolated to a microkernel, independent of third-party drivers. Host portions of the Hyper-V activities are isolated in a parent partition, separate from each guest. The parent partition itself is a virtual machine. Each guest virtual machine operates in its own child partition.

These are the recommended security best practices on a Hyper-V environment,

cumulative to the usual security best practices for physical servers:

- Consider using Domain Isolation with IPSec for both Hosts and Guests.
- Securing the communications between the Hyper-V server and its administrators and users.

Host Operating System Configuration

- Use a Server Core installation for the management operating system.
- Keep the management operating system up to date with the latest security updates.
- Use a separate network with a dedicated network adapter for the management operating system of the physical Hyper-V computer.
- Secure the storage devices where you keep virtual machine resource files.
- Harden the management operating system using the baseline security setting recommendations described in the Windows Server 2008 Security Compliance Management Toolkit.
- Configure any real-time scanning antivirus software components installed on the management operating system to exclude Hyper-V resources.
- Do not use the management operating system to run applications.
- Do not grant virtual machine administrators permission on the management operating system.
- Use the security level of your virtual machines to determine the security level of your management operating system.
- Use Windows® BitLocker™ Drive Encryption to protect resources. (Note: BitLocker does not work with Failover Clustering.)

Virtual Machine Configuration

- Configure virtual machines to use fixed-sized virtual hard disks.
- Store virtual hard disks and snapshot files in a secure location.
- Decide how much memory to assign to a virtual machine.
- Impose limits on processor usage.
- Configure the virtual network adapters of each virtual machine to connect to the correct type of virtual network to isolate network traffic as required.
- Configure only required storage devices for a virtual machine.
- Harden the operating system running in each virtual machine according to the server role it performs using the baseline security setting recommendations described in the Windows Server 2008 Security Compliance Management Toolkit.
- Configure antivirus, firewall, and intrusion-detection software within virtual machines as appropriate based on server role.
- Ensure that virtual machines have all the latest security updates before

they are turned on in a production environment.

- Ensure that your virtual machines have integration services installed.

Network configuration

The Hyper-V server should have at minimum two physical network interface cards (NICs), and potentially more to isolate groups of guest virtual machines from one another.

The first NIC should be used to manage the host partition and the remaining NICs would be used by the guest machines for communication with the physical network and storage. Using separate interfaces is particularly useful because should the NIC(s) in use by child partitions become overloaded, the administrator can still access the host partition.

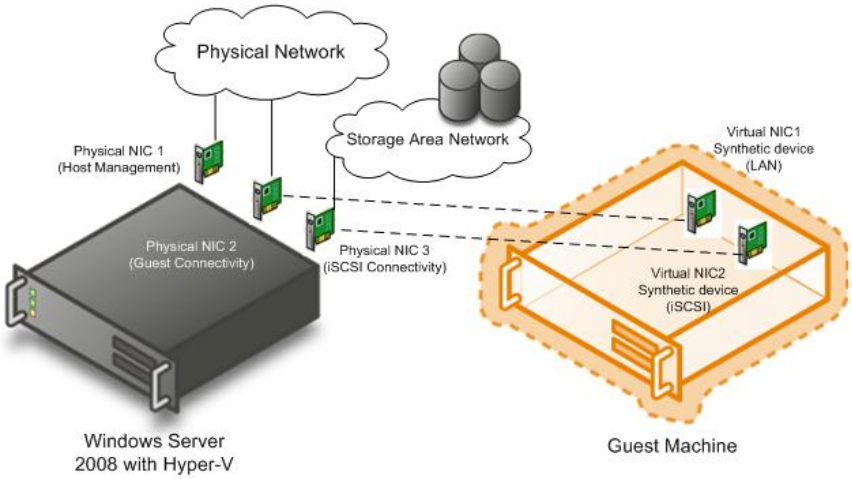

In addition, guest machines with particularly sensitive data could be configured to use only one NIC to access the physical network. With VLANs and other physical boundaries controlling who has access to those systems, administrators can add another layer of security, depending upon access either to an additional physical NIC or a Virtual Network.

Domain Isolation

There are advantages, and relatively small overhead in terms of additional traffic, to implementing IPSec-based domain isolation, especially when utilizing Kerberos-based authentication, in the domain to which the Hyper-V host is joined. Administrators will be assured that only systems that have been authenticated by Kerberos can browse or attach to the Hyper-V host. Domain isolation also blocks the rogue machine plugging into the internal network from browsing and scanning for servers. The intruder will simply get a blank reading as it attempts to list servers; no server will accept its queries.

Since domain isolation utilizes only IPSec authentication to isolate the systems,

the impact on the overall performance of the system is minimal. Unlike the server isolation scenario, IPSec does not encrypt the actual data.

In general, it is recommended to use domain isolation whenever possible inside the virtual environment, and utilize server isolation only where absolutely essential. If there is no way to physically isolate the management console from the rest of the network, server isolation can be used with an IPSec policy to link only the Administrator's console to the management NIC which has access to the Parent partition and full management of the Hyper-V Host.

> **Performance impact**
>
> IPSec hardware accelerators are not effective in virtual environments, and thus cannot help offload IPSec traffic onto hardware.

Recommended Firewall Exceptions for Hyper-V

Below are the required ports that must be open for Hyper-V to operate properly; these are established automatically when the Hyper-V role is added to the Windows 2008 R2 server. They should not be changed either via Group Policy or locally. This configuration can be permanently set in the security policy to make sure other policies do not override and shut down on the essential Hyper-V services.

These ports were extracted from the Windows Server 2008 Hyper-V Attack Surface Reference.xlsx, a reference of all the files, services and ports affected by the Hyper-V role. The spreadsheet can be downloaded from:

http://download.microsoft.com/download/8/2/9/829bee7b-821b-4c4c-8297-13762aa5c3e4/Windows%20Server%202008%20Hyper-V%20Attack%20Surface%20Reference.xlsx

BitLocker

An attacker could gain physical access to the server and access the server's data on the physical drive, accessing the NTFS partition without authentication simply by inserting a Microsoft Windows Pre-installation Environment (WinPE) CD and booting. If data is not encrypted with Encrypted File System (EFS) or some other method, all files will be exposed.

The best response to this is to secure the volumes storing Hyper-V system files and virtual machines with Windows® BitLocker™ Drive Encryption, a hardware based volume encryption which is built into Windows Server 2008.

> **Performance impact**
>
> Volume encryption with any technology adds a small overhead to the server. There is no official document on this subject, but testing by the Product Group shows worst case 8%, and usually between 3-5% hits on performance once BitLocker is turned on. Test performance metrics before and after adding BitLocker and enabling volume encryption.

Administrative Rights Delegation

When a single physical server is configured to support multiple operating system instances, the question of who is granted administrative privileges to which instances becomes important in the context of securing the Hyper-V environment.

Authorization Manager (Azman.msc) is part of the Windows Role-Based Access Control (RBAC) Framework. It is used to delegate administrative rights so that users can perform designated tasks (operations) based on role definitions and assignments. The default scope allows only members of the administrators group the right to create and control virtual machines.

> **Note**
>
> If Microsoft® System Center Virtual Machine Manager is being used, any Authorization needs to be configured from within the Virtual Machine Manager console rather than using AzMan.

These are the main AzMan concepts:

- **Scope**: A collection of similar resources which will share the same authorization policy, for instance, a virtual machine or a virtual network.
- **Role**: A job category or responsibility. Examples: Administrators; Self-Service Users (in Virtual Machine Manager)
- **Task**: A collection of operations or other tasks. Examples: Manage Hyper-V server settings, Create virtual machines.
    - **Operation**: Operations are sub-components of tasks, or can be assigned to a role individually. An operation is an action that a user can perform. Examples: “Start virtual machine”; “Stop virtual machine”. Grouping operations creates a task, and the task permits the role to perform specific administrative functions.

HOST SIZING AND CONSOLIDATION PLANNING

Host sizing consists of determining the total of all the workloads to be consolidated (CPU, RAM, Disk I/O, Network I/O, and so on) as well as the largest individual workloads to be consolidated. Next, a standard host server architecture (or architectures) is designed and benchmarked to determine the real world capacity of each host server architecture. The capacity of the individual standard host server is divided into the workload to be consolidated to determine the number of host servers required. This is done individually for processor, RAM, Disk I/O, Network I/O, and so on. This number is used to determine the overall consolidation ratio by comparing it to the original number of candidate servers.

The end result of the sizing methodology is one or more standard host server architectures and the required number of servers required of each architecture in order to virtualize the complete candidate workload.

In most cases, the amount of RAM in the host server is the factor that first constrains the number of guests the server is able to host. Additionally, the amount of RAM allocated to an individual guest is usually the constraining factor for guest OS performance. Fortunately, the price of RAM and the maximum amount of RAM supported by commodity servers has improved significantly in the last few years. As such, this guidance recommends host servers with a large amount of RAM and the allocation of at least 2 GB of RAM for production guest virtual machines. While it is certainly possible to achieve much higher consolidation ratios, the previous guidance will ensure high performing virtual machines.

Consolidation Candidate Workload Analysis

During the Discovery and Assessment phases, a set of physical servers that are consolidation candidates is identified. The consolidation candidates are analyzed over a period of time to determine the average and peak utilization of CPU, RAM, Disk I/O, Network I/O, and so on. This workload analysis is a critical part of the host sizing process because it determines the total workload that must be supported on the host architecture as well as the largest individual workloads that are going to be consolidated to ensure that they are within the physical limits of a single virtual machine (4 CPU Core, 64 GB of RAM, and so on). If a single workload exceeds the limits of a single virtual machine, it should not be migrated to a virtual machine without an architecture change to scale it out to additional virtual machines.

Host Server Architecture Pattern

When determining the capacity available for guests, reserve one CPU core, 1 GB of RAM, 1 NIC, and a RAID 1 partition for the host server itself sized at 20 GB + the amount of RAM in the server. For example, if the server has 32 GB of RAM, the RAID 1 OS partition would be sized to a minimum of 52 GB to ensure that the proper amount of space required for memory dumps in the case of host OS failures. The remainder of the host server capacity can be utilized for guest virtual machines.

Define Guest Hardware Profiles

System Center Virtual Machine Manager provides the concept of Hardware Profiles, which are user defined collections of virtual machine guest hardware settings such as the number of logical processors, the amount of RAM, and so on. When creating a new virtual machine, the hardware profile can be selected to ensure that a consistent set of configuration settings are applied for all guest virtual machines that use the profile.

For simplicity and consistency of consolidation planning, three guest hardware profiles are recommended and detailed below.

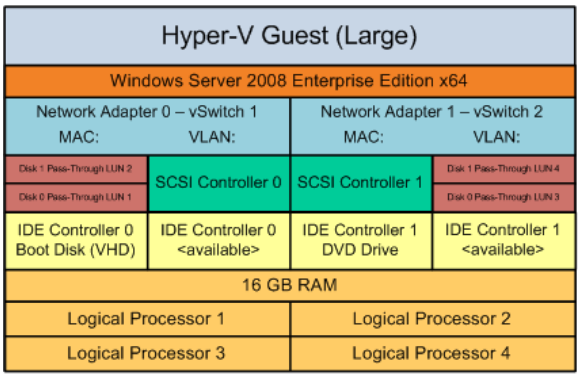

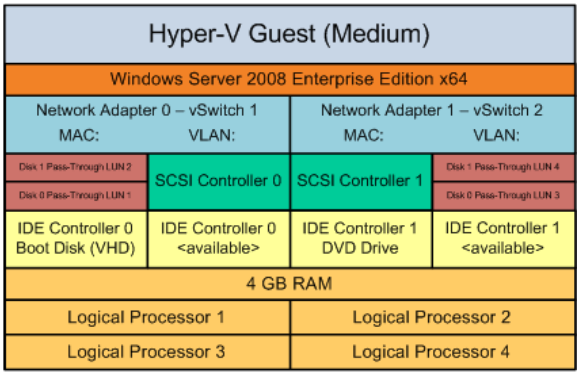

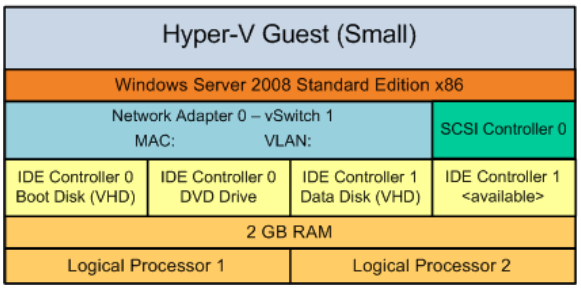

It is recommended that you associate each consolidation with one of the three hardware profiles. The hardware profiles can be used to calculate how many of each profile the host servers can handle by dividing the total CPU, RAM, and I/O capacity of the hosts by the desired mix of guest hardware profiles and the quantity of each.

Benchmark the Host and Guest Server Architecture

Benchmarking of the server's maximum sustained CPU, RAM, Disk I/O, and Network I/O should be performed. Benchmarking provides much more accurate numbers to utilize in the sizing formula while also providing early warning of any serious component or configuration problems. If the performance benchmarks do not meet expectations, careful review of all hardware, software, and configuration should be performed.

Benchmarking of a standard guest or set of guests can be utilized to ensure that the actual performance of the virtual machine guests matches the number calculated using the formula in the next section.

By properly benchmarking the host and guest combinations, you can be confident that the performance of the virtualized infrastructure will meet or exceed expectations. Failure to properly benchmark and confirm the calculations and sizing methodology utilized can result in poor performance.

Calculate the Number of Host Servers Required

Consolidation Candidate Total Resource Requirements

Utilizing the performance data collected during the previous phases, determine the overall resource requirements for all of the consolidation candidates per site.

$$Total\ CPU\ Requirement = \sum_{Candidate\ 1}^{Candidate\ n} \#\ \text{Cores} * \text{CPU Speed (MHz)} * \text{CPU Util (\%)}$$

$$Total\ RAM\ Requirement = \sum_{Candidate\ 1 Candidate\ n} (\text{Total RAM (MB)} * \text{RAM Util (\%)}) + 32\ \text{MB}$$

$$Total\ Network\ I/O\ Requirement = \sum_{Candidate\ 1}^{Candidate\ n} \sum_{\text{NIC 1}}^{\text{NIC n}} \text{Avg Bytes Received per sec} + \text{Avg Bytes Send per sec}$$

$$Total\ Disk\ I/O\ Requirement = \sum_{Candidate\ 1}^{Candidate\ n} \text{Avg Disk Reads/sec} + \text{Avg Disk Writes/sec}$$

Host Server Resources

For each site, the host server architecture pattern and host server architecture should be selected. Using the equations below, divide the host server resources by the consolidation candidate resource requirements. Alternatively, divide the host server resources by the hardware profile resources.

$$Host\ Server\ CPU\ Capacity\ for\ Guests = ((\#\ \text{CPUs} * \#\ Cores\ per\ CPU) - 1) * 85\%$$

$$Host\ Server\ RAM\ Capacity\ for\ Guests = Host\ Server\ Total\ RAM\ (GB) - 2\ GB$$

$$Host\ Server\ Disk\ IO\ Capacity\ for\ Guests = Host\ IO\ Benchmark\ IOPS * .85$$

$$Host\ Server\ Network\ IO\ Capacity\ for\ Guests = Host\ IO\ Benchmark\ Avg\ Bytes\ per\ second * .85$$

Utilizing the above calculations, whichever equation results in the largest number of host servers should be used as the determining factor for the required number of servers for that site.

SYSTEM CENTER VIRTUAL MACHINE MANAGER 2008 R2

System Center Virtual Machine Manager Components

This section provides a brief overview of the System Center Virtual Machine Manager components and some information about each component that should be considered before installation.

System Center Virtual Machine Manager Server

The System Center Virtual Machine Manager server is the hub of a System Center Virtual Machine Manager deployment, through which all the other System Center Virtual Machine Manager components interact and communicate.

The System Center Virtual Machine Manager server runs the System Center Virtual Machine Manager service, which runs commands, transfers files, and controls communications with other System Center Virtual Machine Manager components and with all machine hosts and System Center Virtual Machine Manager library servers. These are collectively referred to as managed computers. The System Center Virtual Machine Manager service is run through the System Center Virtual Machine Manager agents that are installed on the managed computers.

The System Center Virtual Machine Manager server also connects to a Microsoft SQL Server® 2005 database that stores all the System Center Virtual Machine Manager configuration information.

By default, the System Center Virtual Machine Manager server is also a library server, which can be used to store file-based resources such as virtual hard disks, virtual floppy disks, templates, PowerShell™ scripts, unattended answer files , ISO images and System Center Virtual Machine Manager meta data such as hardware

and guest operating system (OS) profiles.

System Center Virtual Machine Manager Administration Console
The System Center Virtual Machine Manager Administrator Console is used to:

- Create, deploy, and manage virtual machines and templates
- Monitor and manage hosts (Windows Server® 2008/ Windows Server® 2008R2 Hyper-V™, Microsoft® Virtual Server 2005 and VMware® Virtual Center managed ESX servers) and library servers
- Manage library objects and jobs
- Manage global configuration settings

The System Center Virtual Machine Manager console is installed after the System Center Virtual Machine Manager server. It can be installed on the same computer as the System Center Virtual Machine Manager server or on a different computer. All the functions that are available through the System Center Virtual Machine Manager Administrator Console are also available by using cmdlets in Windows PowerShell.

System Center Virtual Machine Manager Self-Service Portal v1
The System Center Virtual Machine Manager Self-Service Portal is an optional Web-based component that can be installed and configured to enable end users to create and manage their own virtual machines within a controlled environment.

> **Important**
>
> VMMSSP 2.0 is not an upgrade to the existing VMM 2008 R2 self-service portal. You can choose to deploy and use one or both self-service portals depending on requirements.

System Center Virtual Machine Manager Agent
The System Center Virtual Machine Manager Agent manages virtual machines on virtual machine hosts and allows hosts and library servers to communicate with and transfer files to or from the System Center Virtual Machine Manager server.

When a virtual machine host or library server is joined to a trusted domain and added using the System Center Virtual Machine Manager Administrator Console, System Center Virtual Machine Manager automatically installs an agent on that managed computer by using the default settings.

If a host is on a perimeter network or is not joined to a trusted domain, the agent must be manually installed on the host before it can be added to System Center

Virtual Machine Manager.

Virtual Machine Host

A virtual machine host is a physical computer that hosts one or more virtual machines. Hosts are added to System Center Virtual Machine Manager by using the Add Hosts Wizard in the System Center Virtual Machine Manager Administrator Console. When a host is added to System Center Virtual Machine Manager, an agent is automatically installed on the host system. When you add a Windows® based host, System Center Virtual Machine Manager automatically installs or upgrades the appropriate version of Virtual Server or enables Hyper-V role.

> **Important**
>
> In order to manage Virtual Server hosts running the Windows Server® 2003 operating system the correct version of Windows Remote Management (WinRM) must be installed.

Host Groups

Virtual machine hosts can be organized into groups, which provide ease of monitoring and management of hosts and virtual machines. The host groups can be used in a manner that works best for your organization.

A host group's most basic function is to act as a container to group hosts and virtual machines in a meaningful way. Host groups can be used to:

- Set aside resources on the hosts for the use of the host operating system.
- Designate hosts that are used for self-service.
- Designate which hosts are connected to a storage area network (SAN). (This is a best practice.)
- Enable the automatic placement of virtual machines on the best host in a group of hosts.

Inheritance of Host Group Properties

A child host group can inherit host reserve settings and role delegation from its parent group. However, property inheritance works differently for the following two features:

- **Host reserves**. When changing the host reserves for a parent host group, the administrator can choose whether to cascade the host reserve settings to the hosts in the child groups. If inheritance is enabled, the host reserve settings in the child groups will be overwritten.
- **Role Delegation**. If a parent host group is used for virtual Role

Delegation, each of its child host groups will automatically inherit these settings from the parent

Group Isolation

A host group can be used to isolate a host or collection of hosts. If, for example, a host has virtual guests that host mission-critical applications, that host can be isolated by placing it in its own host group. In this manner, the administrator can be sure that only appropriate users are delegated permissions and that host reserve resources are maximized for availability.

System Center Virtual Machine Manager Library Server

Each System Center Virtual Machine Manager library server contains a catalog of resources that can be used to create and configure virtual machines in System Center Virtual Machine Manager. The library contains files that are stored on library shares, and it can contain file-based resources such as virtual hard disks, virtual floppy disks, ISO images, and scripts.

> **Important**
>
> After the setup is complete, the default library server and library share cannot be moved. Give careful consideration to its location when running Setup.

In addition, the library server can contain virtual machine templates, hardware profiles, and guest operating-system profiles, which can be used to create virtual machines and store virtual machines that are not in use.

Library Groups

As library servers are created, library groups can be created to help organize the library servers in a way to best suit your needs.

It is a best practice to align the library servers with the host groups that use those resources, especially when the library server is SAN-connected. In this way, it is known which hosts and library servers can be used to take advantage of SAN file transfers.

System Center Virtual Machine Manager Server Placement

For most System Center Virtual Machine Manager deployments, a single System Center Virtual Machine Manager server is sufficient. System Center Virtual Machine Manager deployments can be subsequently scaled by adding more virtual machine hosts and library servers as the environment grows. Having a single System Center Virtual Machine Manager server with a single database also

lends itself to the central management of the entire virtual environment. However, having more than one System Center Virtual Machine Manager server can be beneficial in the following situations:

- When the development and test environments are managed separately from the production virtual environment
- When the virtual environments grows—or is planned to grow—beyond the supported maximum of 400 hosts and 8,000 virtual machines

If business needs dictate that more than one System Center Virtual Machine Manager server should be installed, the following points should be considered:

- Each System Center Virtual Machine Manager server must be installed on a separate computer and have a separate System Center Virtual Machine Manager database.
- The movement of files from one System Center Virtual Machine Manager deployment to another is not supported.
-

Supported Number of Hosts and Virtual Machines

The maximum number of hosts and virtual machines tested with and supported by System Center Virtual Machine Manager on the largest recommended hardware configuration is 400 hosts and 8,000 virtual machines. These are practical rather than hard-coded limitations; this number may be revised up or down depending on customer or fault tolerance requirements.

The number of virtual machines that can be run on a host is primarily limited by the configuration of the host and virtual machines.

Network Considerations

From a networking perspective, the key things to consider for System Center Virtual Machine Manager are:

- Connectivity
- Bandwidth
- Network traffic

Connectivity

Ensure that any firewalls that exist do not block the necessary communications among System Center Virtual Machine Manager components.

When System Center Virtual Machine Manager is installed, the ports that the System Center Virtual Machine Manager server uses for agent communication and file transfer between hosts and library servers are specified. By default, those

ports are 22 (SFTP) 80 and 443, respectively.

Bandwidth

Using System Center Virtual Machine Manager to create and manage virtual machines can involve moving multi-gigabyte files across the network—for example, when performing P2V migrations or migrating a virtual machine from one host to another or deploying a new VM from a template.

As a best practice, connect all the computers in a System Center Virtual Machine Manager configuration with at least a 100-MB full duplex Ethernet connection. When using a Gigabit Ethernet connection, more powerful processors than those that are recommended will further improve performance.

When extending System Center Virtual Machine Manager beyond the data center, such as in remote or branch office scenarios:

- Consider adding System Center Virtual Machine Manager library servers to remote locations where there will be a need to provision virtual machines or templates or access ISO images.
- Avoid performing file transfers across small or unreliable WAN links.

Network Traffic

System Center Virtual Machine Manager performs periodic refreshes of the library, hosts, and virtual machines. In very large virtual environments, the amount of traffic can become significant.

If using a Fiber Channel or Internet Small Computer System Interface (iSCSI) SAN, the network impact can be reduced by doing SAN transfers in lieu of network transfers. When a SAN transfer is performed, the LUN that contains the virtual machine is remapped from the source to the destination computer (instead of a network file transfer being performed). Because of this, SAN transfers are much faster and are independent of the size of the files being transferred. When using iSCSI, consider the network traffic that will be incurred by using iSCSI connections with System Center Virtual Machine Manager.

Storage Considerations

System Center Virtual Machine Manager supports all forms of direct-attached storage (DAS) as well as Fiber Channel and iSCSI SANs. System Center Virtual Machine Manager also supports N_Port ID Virtualization (NPIV) on a Fiber Channel SAN. NPIV makes use of the host bus adapter (HBA) technology that creates virtual HBA ports on hosts by abstracting the underlying physical port. This enables a single HBA port to function as multiple logical ports, each with its own identity. Each virtual machine can then attach to its own virtual HBA port

and be independently zoned to a distinct worldwide name (WWN).

SAN Transfers with System Center Virtual Machine Manager

System Center Virtual Machine Manager can perform the following types of SAN transfers between a source and a destination computer:

- Storing a virtual machine from a virtual machine host in a System Center Virtual Machine Manager library
- Deploying virtual machines from a System Center Virtual Machine Manager library to a host
- Migrating a virtual machine from one host to another

When a SAN transfer is performed, the LUN that contains the virtual machine is remapped from the source to the destination computer instead of the files being transferred over the network. Therefore, SAN transfers are much faster than standard network transfers and are independent of the file size.

If SAN transferring is available, System Center Virtual Machine Manager will use it automatically. This behavior can be overridden to force System Center Virtual Machine Manager to use a network transfer.

Before System Center Virtual Machine Manager can be used for a SAN file transfer, the following configuration steps must be completed:

1. Install Virtual Disk Service (VDS) 1.1, a component of Windows Server 2003 R2, on each computer that will serve as either a source or destination.
2. Install the VDS hardware provider only on the System Center Virtual Machine Manager server.
3. Install an iSCSI initiator for an iSCSI SAN.
4. Install a multipath I/O (MPIO) driver for a Fiber Channel SAN, even if it is using only one HBA port.

Before System Center Virtual Machine Manager can be used for a SAN file transfer on Windows Server 2008/Hyper-V Hosts, the following configuration steps must be completed:

1. Install an iSCSI initiator for an iSCSI SAN.
2. Install a multipath I/O (MPIO) driver for a Fiber Channel SAN, even if it is using only one HBA port.

System Center Virtual Machine Manager 2008 R2 Rapid Provisioning with SAN

Some SANs have resources to clone a LUN containing a VHD and presenting it to the host. In order to use System Center Virtual Machine Manager for the OS customization and IC installation, System Center Virtual Machine Manager R2

provides the switch UseLocalVirtualHardDisk for the new-VM cmdlet without the network copy. You can create a template which includes the OS answer file and references a dummy VHD which is not used. This feature is only available using Windows PowerShell™.)

This is a sample script:

```
Get-VMMServer -ComputerName "VMMServer1.Contoso.com"
$JobGroupID = [Guid]::NewGuid().ToString()
$Template = Get-Template | where {$_.Name -eq MyTemplate"}
$VMHost = Get-VMHost | where {$_.Name -eq "VMHost.Contoso.com"}
Move-VirtualHardDisk -IDE -BUS 0 -LUN 0 -Path "L:\OS.VHD" -JobGroup
$JobGroupID
New-VM -Name "VM Name" -Path "L:\" -Template $Template -VMHost $VMHost -
JobGroup -$JobGroupID -UseLocalVirtualHardDisks
```

Security Considerations

General Security Considerations

Consider the following information when planning a System Center Virtual Machine Manager deployment:

- When using multiple Active Directory forests, a bidirectional trust relationship is required to install System Center Virtual Machine Manager components.
- By default, virtual machines are run in the security context of the account that started the machine. For enhanced security, an account with a low level of privileges can be specified.
- When using a remote instance of SQL Server, the instance must run under an account other than LocalSystem.
- Self-service users might be asked for credentials when connecting to virtual machines. To avoid this, add the host name to the Local Intranet sites in Security Settings or Microsoft® Internet Explorer®.
- When adding a virtual machine host or library server, System Center Virtual Machine Manager installs the System Center Virtual Machine Manager server's machine account as a local Administrator on the managed computer. Ensure that groups restricted by Group Policy do not remove this account; otherwise, System Center Virtual Machine Manager will not function correctly.

Security Vulnerabilities

To avoid common security vulnerabilities, consider the following:

- As a best practice and where practical, avoid using the default ports when installing System Center Virtual Machine Manager components.
- Firewall and antivirus software that are running on the host operating system do not protect guest virtual machines. For optimal production, install these products on the guest operating systems in addition to the host.
- Limit access to the host file system. The access control list (ACL) for library shares should contain only System Center Virtual Machine Manager Administrators, the System Center Virtual Machine Manager server's machine account, and self-service users (if appropriate).
- When a virtual machine host or library server is added, System Center Virtual Machine Manager remotely installs a System Center Virtual Machine Manager agent on the managed computer. This process opens a range of DCOM ports and uses Server Message Block (SMB). If this is a concern for the customer, the System Center Virtual Machine Manager agent can be manually installed on the host and then remotely discovered from the System Center Virtual Machine Manager Administrator Console by using only the Microsoft Windows® Remote Management (WinRM) port (80 by default) and the Background Intelligent Transfer Service (BITS) port (443 by default).
- To create and manage virtual machines on a host, an administrator needs to have been assigned the appropriate role and does not require local administrative privileges.

Monitoring and Reporting

Reporting in System Center Virtual Machine Manager is provided though the Server Virtualization Management Pack for System Center Operations Manager 2007. Before reports can be viewed and used, Operations Manager must be installed and the Server Virtualization Management Pack deployed. Reports are generated by Operations Manager but can be opened in Reporting view in the System Center Virtual Machine Manager Administrator Console.

Additionally, the Operations Manager 2007 agent must be installed on each machine that will be monitored.

One report that is helpful in planning the virtual environment is the Virtualization Candidates report. This report helps to identify the physical computers that are good candidates for conversion to virtual machines. The Virtualization Candidates report can be used to identify little-used servers and to display the average values for a common set of performance counters for CPU, memory, and disk usage along with hardware configuration information, including processor speed, number of processors, and RAM. The report can be limited to computers

that meet the specified CPU and RAM requirements, and it can sort the results.

The Server Virtualization Management Pack discovers the following objects:

- System Center Virtual Machine Manager Managed Virtual Machine
- System Center Virtual Machine Manager Agent
- System Center Virtual Machine Manager Managed Host
- System Center Virtual Machine Manager Host Group
- System Center Virtual Machine Manager Engine Server
- System Center Virtual Machine Manager Library Server
- System Center Virtual Machine Manager Database
- System Center Virtual Machine Manager Self-Service Server
- System Center Virtual Machine Manager Self-Service Website
- System Center Virtual Machine Manager Management Group
- Virtual Server 2005 R2
- Virtual Machine
- Virtual Machine Computer

Planning for Physical-to-Virtual Migrations

Where supported by the candidate operating system, System Center Virtual Machine Manager will be the primary method by which physical workloads are migrated to virtual machines. In System Center Virtual Machine Manager, a P2V conversion is the process by which a functioning physical computer is copied to an identical—or nearly identical—virtual machine. During a P2V conversion, disk images of the physical hard disks of the target computer are created and formatted as virtual hard disks (.vhd files) for use in the new, virtual machine. The new, virtual machine will have the same identity as the original, physical machine upon which it is based.

System Center Virtual Machine Manager can perform either an offline or an online P2V migration on all supported operating systems (only offline is available for Windows 2000). Online conversions rely upon the Volume Shadow Copy Service (VSS), so the source machine does not have to be rebooted during the process. In an offline conversion, the source machine is rebooted into the Windows® Pre-installation Environment (Windows PE) to image the physical disks.

The following table lists the online and offline support for P2V migrations using System Center Virtual Machine Manager 2008 R2:

Operating system	P2V	P2V	V2V

	offline	online	
Windows Server® 2008 / Windows Server® 2008 R2 with Hyper-V role enabled	No	No	No
Windows Server® 2008 / Windows Server® 2008 R2 without Hyper-V role enabled	Yes	Yes	Yes
Windows Server® 2003 SP1 or later	Yes	Yes	Yes
Windows Server® 2003 x64 Edition	Yes	Yes	Yes
Windows® 2000 Server SP4	Yes	No	Yes
Windows® XP SP2 or later	Yes	Yes	Yes
Windows® XP x64 Edition	Yes	Yes	Yes
Windows Vista®	Yes	Yes	Yes
Windows Vista® x64	Yes	Yes	Yes
Windows® 7	Yes	Yes	Yes
Windows® 7 x64	Yes	Yes	Yes

If there are consolidation candidates that are running Microsoft Windows NT® 4.0 or other operating systems or service packs that are not supported by System Center Virtual Machine Manager you can use the Microsoft Virtual Server 2005 Migration Toolkit (VSMT) or third-party tools as spot solutions for these servers.

Migration Requirements

Before beginning a P2V conversion with System Center Virtual Machine Manager, review the following requirements and limitations:

- **Host server requirements**. A P2V conversion requires that the target virtual machine host is a computer running Hyper-V or Virtual Server 2005 R2 SP1.
- **Online versus offline P2V conversions**. In an offline P2V conversion, the source machine is booted into Windows PE to image the physical disks. In an online conversion, VSS is used, and the machine does not have to be rebooted prior to migration.
- **Offline P2V memory requirement**. An offline P2V conversion requires

the source machine to have at least 512 MB of physical memory. Additional storage and network drivers may need to be provided so WINPE can successfully image the source machine.

- **Updates requirement (if needed)**. A P2V conversion may require that additional files be added to the internal System Center Virtual Machine Manager patch cache. In this case:
    - Use the information that is provided by the wizard to identify which updates are required.
    - Obtain the patch files and copy them to the **Patch Import** directory on the System Center Virtual Machine Manager server.
    - Click **Check Again** to continue running the wizard.
- **Bad sectors that do not transfer**. Bad sectors on a disk cannot be transferred during a P2V conversion. To avoid data loss, run a disk maintenance tool (such as chkdsk) on the source machine prior to migration.

Pre-Migration Considerations

To be assured of the maximum chance of success in performing P2V migrations, additional considerations must be taken into account. The following sections describe these considerations.

Testing

The methodology for performing the P2V migration should be thoroughly tested and documented in an isolated lab environment prior to performing the migration in the production environment. Ensure that the following areas of migration are thoroughly tested and documented:

- **Pre-migration server preparation**. Test and document the procedures for ensuring that the P2V candidate machine is in good working order by running chkdsk and defrag and performing the applicable hardware diagnostics.
- **System Center Virtual Machine Manager configuration**. Verify that the recommended configuration and placement of System Center Virtual Machine Manager suits your needs.
- **Disaster recovery**. Prepare contingency plans for recovering failed P2V migrations or for restoring service to the source hardware in the event of unplanned complications post-migration.
- **P2V tool usage**. Some P2V scenarios will include operating systems that are not directly supported by System Center Virtual Machine Manager, necessitating the use of VSMT or a third-party tool for the P2V migration.

Migration Ordering

Start with locations within the data center and the P2V candidates that afford the best opportunity for success. This approach will build confidence and proficiency in the virtualization team prior to tackling remote sites; branch offices; and other, more challenging scenarios.

When considering remote sites and branch offices, verify that the required hardware and server roles are in place and functional prior to deploying the virtualization team to that location. Ensure that sufficient time is built into the project plan to allow the team members to travel to the remote sites.

Business Continuity

Although P2V migrations are not data-destructive to the source machine, care must be taken to ensure that the application hosted by the virtualization candidate can be recovered in the event of an issue during migration. Ensure that the virtualization team is well aware of and able to perform any applicable disaster recovery steps that are documented, and that full backups of the source machines are taken prior to attempting the migration.

Keep the migrated physical server in place as a cold spare for a period of time after the migration. Generally, two work weeks is sufficient. A longer period of time may be required if there are specific periods of heavy activity that would expose an unanticipated performance issue.

User Acceptance Testing

Prior to placing a migrated virtual machine into production, ensure that personnel are on hand who can perform user acceptance testing to validate the successful migration. In some cases, it may be necessary to perform a P2V migration and leave the source server in service for a short period of time after migration until user acceptance testing (UAT) can be performed.

For complex, critical, or time sensitive applications, it may be necessary to perform the P2V migration and move the migrated virtual machine to an isolated lab environment to perform the necessary validation.

If the virtualization candidate is a domain controller, be mindful of the risk of USN rollback if the server is not to be immediately placed online.

Communications Planning

Where documented, use established maintenance windows to conduct the P2V migrations. Identify all the users of a particular service or application and communicate the planned migration well in advance. For example:

- When virtualizing Active Directory domain controllers, communicate to all the users in that site. Where applicable, verify that the alternate Active Directory site is available and functional and that the site to be migrated

is not an alternate site for another location.

- For line-of-business (LOB) applications, ensure that all the users and stakeholders in the affected organization are communicated to.
- For other services, such as file and print, notify all the users that use files on that particular server.

Ensure that the migration plans are communicated to the personnel in the Network Operations Center and Help Desk so that issues related to migration can be quickly addressed.

SYSTEM CENTER VIRTUAL MACHINE MANAGER SELF SERVICE PORTAL 2.0 (VMMSSP)

VMMSSP Components

This section provides a brief overview of the VMMSSP components and some information about each component.

VMMSSP website
A Web-based component that provides a user interface to the self-service portal. Through the VMMSSP website, users can perform various tasks such as pooling infrastructure assets in the self-service portal, extending virtual machine actions, creating business unit and infrastructure requests, validating and approving requests, and provisioning virtual machines (using the self-service virtual machine provisioning feature). Users can also use the VMMSSP website to view information related to these tasks.

VMMSSP database
An SQL Server database that stores information about configured assets, information related to business units and requests, and information about what has been provisioned to various business units. The database also stores the XML that encodes default and customized virtual machine actions and other information related to the configuration of the self-service portal.

VMMSSP server
A Windows service that runs default and customized virtual machine actions that the user requests through the VMMSSP website. The service uses a Windows Communication Foundation (WCF) TCP endpoint to listen for client communication, and hosts a Windows Workflow Foundation (WF) run-time environment. Using WF, the server component runs the sequences of tasks that

comprise virtual machine actions. You can optimize the performance of the server component using parameters available in the self-service portal or in configuration files; these parameters control or "throttle" the number of operations that can run simultaneously

VMMSSP Reporting Dashboard

The VMMSSP uses a reporting dashboard that is built on Windows SharePoint Services. The Dashboard uses SharePoint Dashboard Configuration and Viewer web parts to supply the reporting from the VMMSSP. There is a default dashboard supplied and there is also facility to create custom reports.

> **Note**
>
> Windows SharePoint Services 3.0 SP2 is the base requirement; however, SharePoint Server 2007 is a supported alternative configuration.

Hardware Requirements

The following table provides the minimum and recommended hardware requirements.

Hardware Component	Minimum	Recommended
RAM	2 GB	4 GB
Available hard disk space	50 GB	50 GB

Software Requirements

Before you install the Self Service Portal components, install and configure the following software on the computer.

Software	Comments
Operating System: Windows Server® 2008 R2	Windows Server 2008 R2 Enterprise Edition and Windows Server 2008 R2 Datacenter Edition are supported.
Windows Server Internet Information Services (IIS) 7.0	You must add the Web server role (IIS) and then install the following role services: • IIS 6 Metabase Compatibility • Static Content • Default Document • ASP.NET

	• .NET Extensibility • ISAPI Extensions • ISAPI Filters • Request Filtering Use Integrated Windows authentication (NTLM or Kerberos). Turn off Anonymous authentication. For more information, see Configure Windows Authentication in the IIS documentation. Use IIS v6.0 compatibility mode.
Microsoft .NET Framework 3.5 SP1	
Windows PowerShell™ 2.0	**Important** If your extensibility scripts require specific Windows PowerShell snap-ins, install them when you install the toolkit server component. **Note** If the Windows PowerShell execution policy is set to **Restricted**, the Setup wizard changes it to **AllSigned**.
Microsoft Message Queuing	Also known as MSMQ.
VMM 2008 R2 Administrator Console	
SQL Server 2008	SQL Server 2008 Enterprise (64-bit) and SQL Server 2008 Standard (64-bit) versions are supported.

VMMSSP Architecture Patterns

Single Server Architecture

The single server architecture pattern is illustrated below. The architecture consists of a single host server running Windows Server 2008 R2(Physical or virtual machine) with SCVMM 2008 R2 administrator console, SQL Server 2008 and SSP components installed.

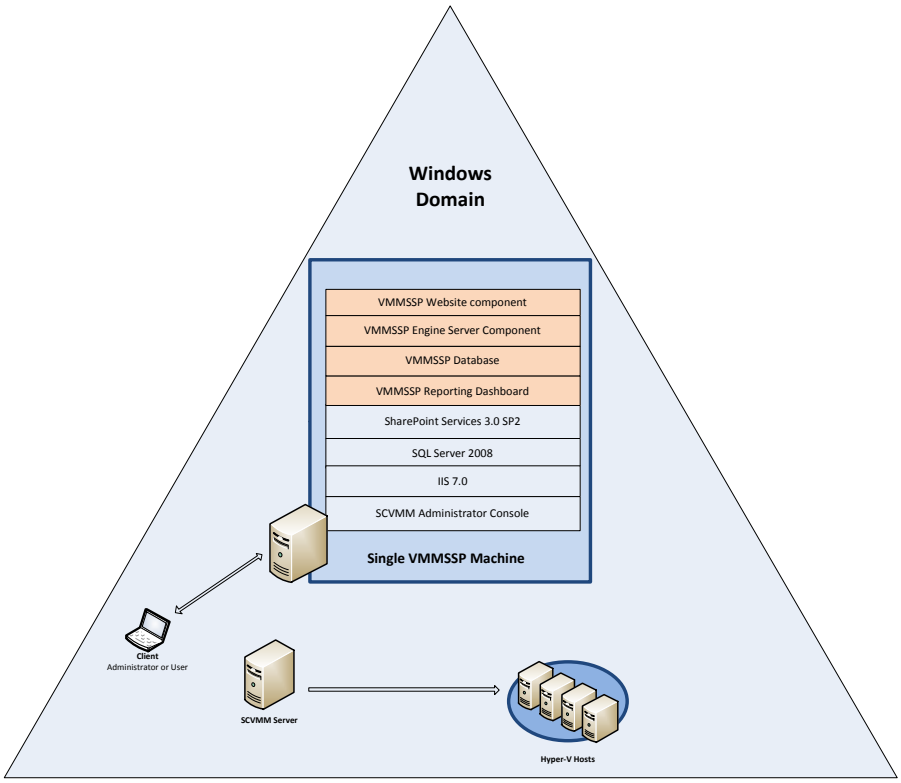

This pattern is appropriate for test and development environments and small branch offices.

Four Server Architecture

The four server architecture pattern is illustrated below. The architecture consists of four servers running Windows Server 2008 R2 (Physical or virtual machine) with SCVMM 2008 R2 administrator console and server component installed on first machine, SQL Server 2008 on the second server, Windows SharePoint Services 3.0 SP2 or Windows SharePoint Server 2007 on the third and SSP web components installed on the fourth .

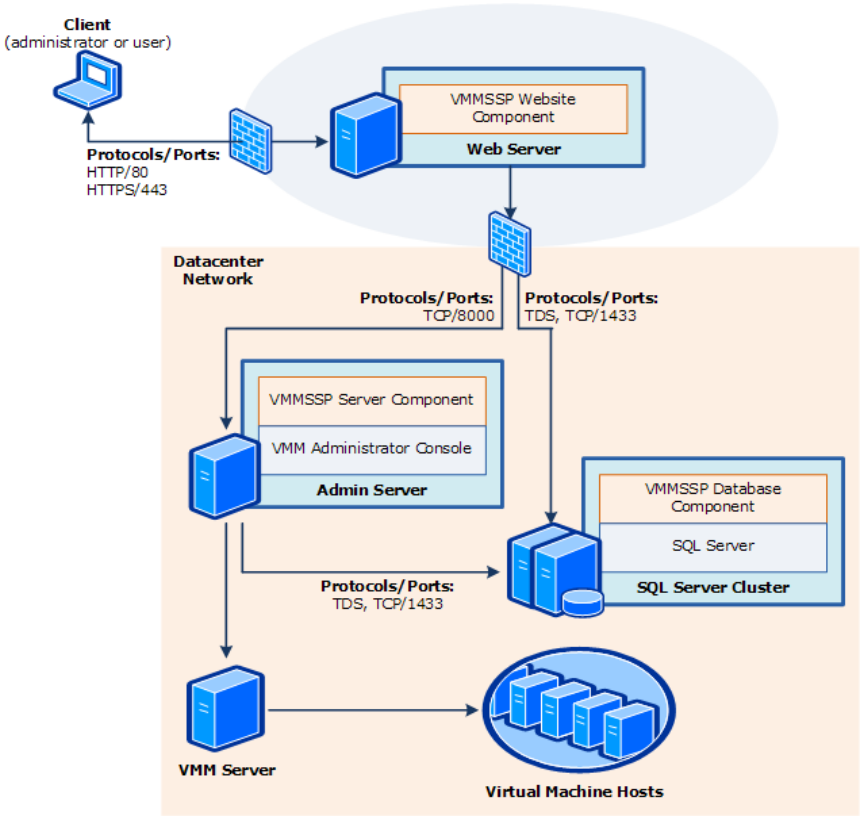

This pattern is appropriate for larger environments where the need to scale out is required.

Security Considerations

Securing the environment for the self-service portal involves the following tasks:

- Understanding and planning the default and custom user roles that are defined in the self-service portal.
- Planning and preparing the service accounts.
- ・Understanding the ports and protocols required for establishing communication channels between various self-service portal components.
- Hardening the Web server that will run the VMMSSP website component.

Accounts and Groups for the SSP User Roles

The self-service portal provides four default user roles; you can also create custom user roles. The self-service portal uses Windows authentication, so you can populate these user roles with Active Directory user accounts and security groups. Plan how you will map security groups to the user roles. In particular, identify the security groups and user accounts that you will add to the DCIT Admin built-in user role. Members of this role are super-administrators; they can perform the entire set of tasks that VMMSSP website exposes. You can add

members to the DCIT Admin role when you run the Setup wizard, as well as adding them later when you configure the self-service portal.

Service Accounts
If you are using a domain account and your domain Group Policy object (GPO) has the default password expiration policy set as required, you will either have to change the passwords on the service accounts according to the schedule, or configure the accounts so that the passwords never expire.

Firewall Exceptions
If Windows Firewall is configured on the computers on which you plan to install the self-service portal, you must ensure that port exceptions are added to the Firewall on those computers that you plan to use for the self-service portal.

To configure any other firewall, refer to the instructions provided by the firewall's manufacturer.

Hardening the Self-Service Portal Website
Installing the VMMSSP website component creates a corresponding website for the self-service portal in IIS. This section specifies the recommendations for hardening the self-service portal website.

Configure SSL for the Self-Service Portal
To encrypt communications between the client and the VMMSSP website component, you should configure SSL security on your Web server. You can obtain the encryption certificate you need for SSL in one of the following ways, depending on how your portal is used:

- If the website is on your organization's intranet, with no public access, you can obtain the certificate from your organization's existing public key infrastructure (PKI).
- If users can access the self-service portal from the Internet, Microsoft recommends that you obtain a certificate from a certification authority.

If you are using IIS 7.0, see Securing Communications with Secure Socket Layer (SSL) in the IIS documentation for more information.

Disabling ISAPI Handlers That Are Not Needed
When you install the VMMSSP website component, IIS lays down the default ISAPI filters and handlers for common extensions such as .soap, .xoml, and .asmx. To avoid unnecessary exposure to any potential security risks, it is recommended that you disable the handlers that the website component is not using.

The Table below lists the ISAPI Handlers for the VMMSSP Website Component.

OPTIONSVerbHandler

PageHandlerFactory-ISAPI-2.0
PageHandlerFactory-ISAPI-2.0-64
TRACEVerbHandler
WebServiceHandlerFactory-ISAPI-2.0
WebServiceHandlerFactory-ISAPI-2.0-64
StaticFile
AXD-ISAPI-2.0
AXD-ISAPI-2.0-64

The following procedure explains how to disable ISAPI handlers in IIS 7.0.

> **Important**
>
> To avoid unintended effects in other Web sites, be careful to only update the handlers for the IIS Web site configured for the self-service portal.

To disable ISAPI handlers for the self-service portal

1. On the Web server, in Administrative Tools, open Internet Information Services (IIS) Manager.
2. Expand Sites, and navigate to the IIS website configured for the self-service portal.
3. In the Features View pane, under IIS, open Handler Mappings.
4. For each handler that is not listed in the preceding table, select the handler, click Remove, and then click Yes.

Monitoring and Reporting

VMMSSP Dashboard

The Microsoft® System Center Virtual Machine Manager Self-Service Portal (VMMSSP) Dashboard is a Windows® SharePoint® Services–based application that provides a view of multiple sets of self-service portal statistics on a single Web page. Users can view data in the form of pie charts, graphs, or Dundas gauges.

The VMMSSP Dashboard supplements the Virtual Machine Manager 2008 R2 Self-Service Portal 2.0 by providing a centralized view of infrastructures, resources, virtual machines, and charge-back data. For each of these areas, the

Dashboard also provides detailed status information. The VMMSSP Dashboard provides the information that IT Managers need to make decisions, to reduce the costs of services, and to improve the overall productivity of the datacenter.

Because the Dashboard is built on Windows SharePoint Services, users can access it without using the self-service portal

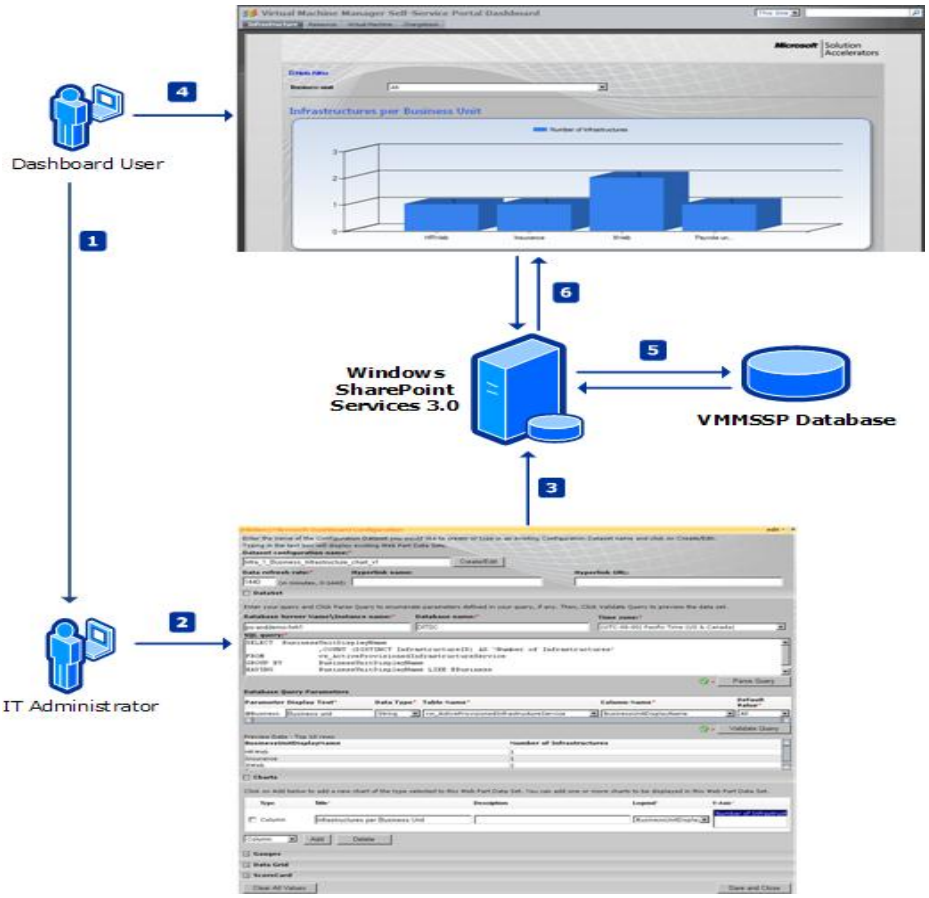

ADDITIONAL RESOURCES

Below are several other resources available to accelerate a successful Server Virtualization deployment.

Microsoft Solution Accelerators

Microsoft provides tools and guidance to help you solve your deployment, planning, and operational IT problems. They are free and fully supported.

Microsoft Assessment and Planning (MAP) Toolkit
Download this network-wide inventory and assessment tool to determine the virtualization candidates for Windows Server 2008 R2 Hyper-V and Application Virtualization. If your customer is currently running VMware, the toolkit now includes a VMware discovery feature that identifies already-virtualized servers running under VMware that can be managed with System Center Virtual Machine Manager or which can be migrated to Hyper-V.

Learn more at: http://technet.microsoft.com/en-us/solutionaccelerators/dd537570.aspx?SA_CE=VIRT-MAP-WEB-SAT-2009-07-13

Offline Virtual Machine Servicing Tool 2.1
The Offline Virtual Machine Servicing Tool 2.1 has free, tested guidance and automated tools to help keep offline virtualized machines updated, without introducing vulnerabilities into your IT infrastructure. The tool combines the Windows Workflow programming model with the Windows PowerShell™ interface to automatically bring groups of virtual machines online, service them with the latest security updates, and return them to an offline state.

Learn more at: http://technet.microsoft.com/en-us/library/cc501231.aspx?SA_CE=OVMST21-Release-VIRTPROD-2009-12-07

Infrastructure Planning and Design Guides for Virtualization
Streamline your virtualization-infrastructure design processes with planning guidance from Infrastructure Planning and Design Guides for

Virtualization. Each guide addresses a unique virtualization-infrastructure technology or scenario, provides critical architectural decisions to be addressed with available options, and supplies the means to validate design decisions to ensure that solutions meet the requirements of both business and IT stakeholders.

Learn more at: http://technet.microsoft.com/en-us/solutionaccelerators/ee395429.aspx

Microsoft.com

In addition to the resources above, please visit http://www.microsoft.com to find resources for delivering Microsoft Server Virtualization technologies.