XenDesktop Design Guide

For Microsoft Windows 2008 R2 Hyper-V
Version 1.2

Citrix Worldwide Consulting Solutions
12 April 2010

www.citrix.com

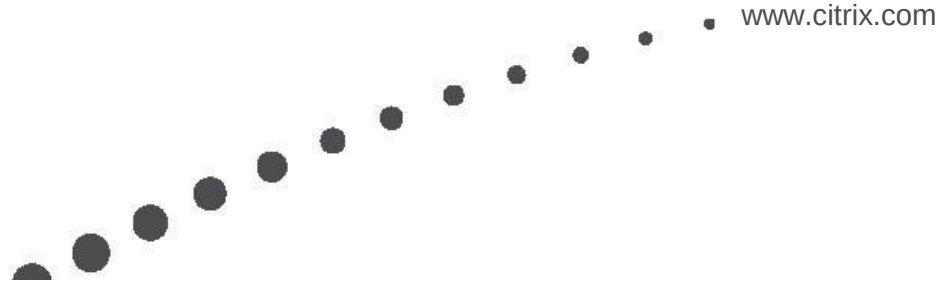

Contents

Introduction 1

XenDesktop Architecture for Hyper-V 1

Infrastructure 2

- Hyper-V 2
  - Choosing the Correct Installation of Hyper-V 2
  - Hardware Considerations 3
  - Desktop High-Availability (Clustering) 4
  - CPU Considerations 5
  - Networking 5
- Virtual Machine Manager 6
  - Design Recommendations 6
  - Integrating with XenDesktop 7
  - Configuring for Disaster Recovery 8
- Storage Selection 8
  - Storage Types 8
  - Saved State Files 9

Operating System Delivery 10

- Operating System Selection 10
- Virtual Hard Disk Considerations 10
  - Avoid Dynamic VHDs 10
  - Manually Configure VHD Partition Offset 11
- Networking Considerations 11
  - Single Network Adapter Configuration 12
  - Dual Network Adapter Configuration 12

Decision Summary Table 13

Deployment Considerations 15

Desktop Delivery Controller Settings 15

XenDesktop Creation 15

Assigning write-cache drives in bulk 16

Conclusion 17

Appendix 18

Configuring Partition Offset Manually with XenConvert 2.1 18

Using BindCfg to Set the Network Card 18

PowerShell Script to Copy and Attach the VHD from VMM 19

Configuring a Dedicated Farm Master 21

Introduction

The purpose of this document is to provide guidance around designing a virtual desktop deployment that leverages Microsoft Hyper-V and Citrix XenDesktop 4 for the infrastructure. This document is not to be considered a replacement for other XenDesktop design guidance, but rather as an addendum that will assist in design decisions specifically related to using Microsoft Hyper-V as the hypervisor. This document represents a more detailed level of guidance from the XenDesktop Modular Reference Architecture available at http://support.citrix.com/article/CTX124087. The information contained in this document was gathered during the joint Citrix and Microsoft testing of XenDesktop on Hyper-V.

This document assumes the reader is already familiar with the basic architecture of XenDesktop and has a detailed understanding of the functionality. This document is not a step-by-step guide for implementing Microsoft technologies such as App-V or System Center with XenDesktop but rather this document focuses on optimizing a Hyper-V installation with XenDesktop based on business requirements.

For simplicity, the document starts with an overview of the XenDesktop architecture when deployed using Microsoft Hyper-V. The next section covers the design decisions that are specific to the Hyper-V environment. The final portion discusses some of the considerations when deploying XenDesktop to a large environment. After reading this document, the reader will be able to effectively design a XenDesktop deployment on Microsoft's hypervisor.

XenDesktop Architecture for Hyper-V

Figure 1 shows the basic design for XenDesktop when using Microsoft Hyper-V as the host hypervisor. Although most of the concepts discussed in this document equally apply to virtualizing servers on Microsoft Hyper-V, the focus of this document is on virtualizing the desktops.

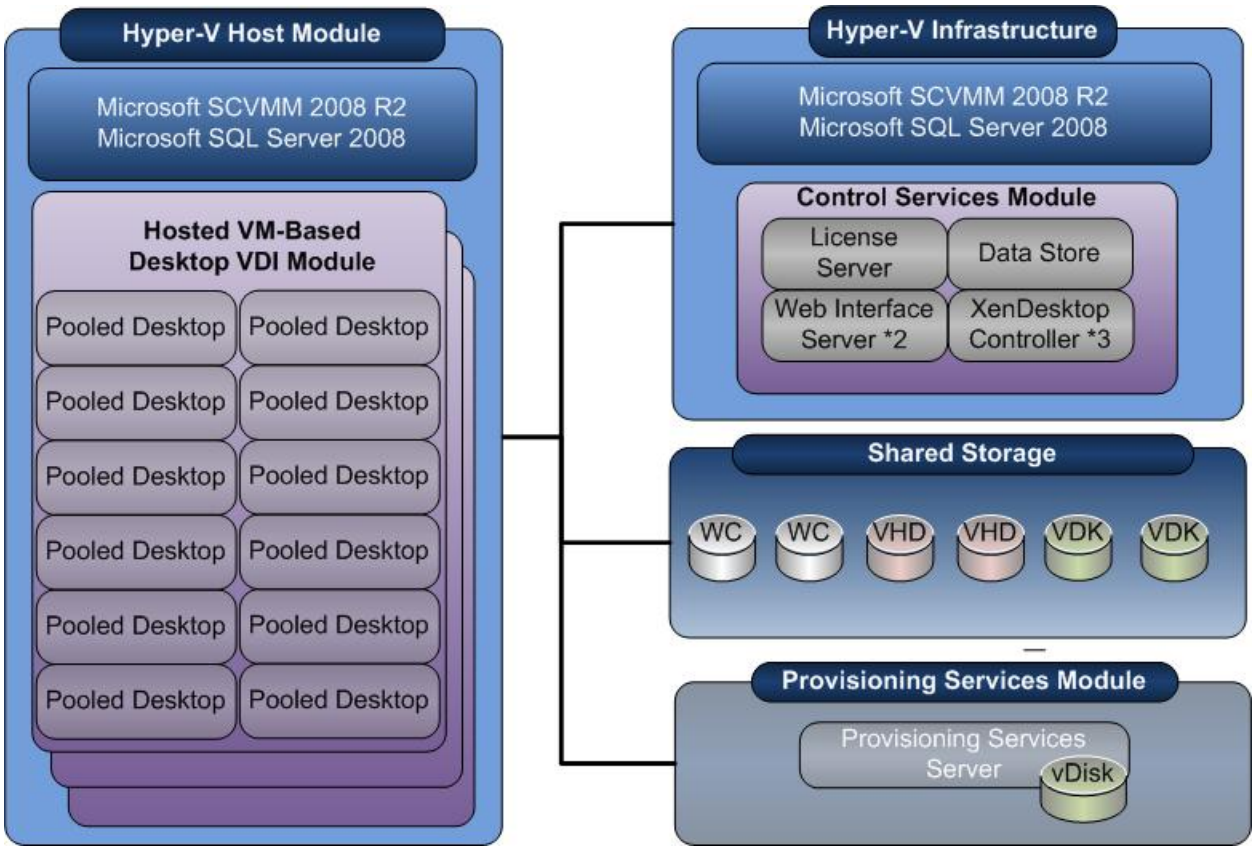

Figure 1: XenDesktop Hyper-V Architecture

Infrastructure

The key design decisions for infrastructure are divided into four sections: Hyper-V, System Center Virtual Machine Manager (VMM), XenDesktop, and Storage. The sections below include architectural guidance and scalability considerations that should be addressed when planning a XenDesktop deployment hosted on Microsoft Hyper-V.

Hyper-V

The Microsoft Hyper-V role brings with it several key decision points. The first of which is what version of Hyper-V is the best for the deployment.

Choosing the Correct Installation of Hyper-V

Using Windows Server 2008 R2 is a must for anyone deploying with the Hyper-V role. Microsoft has added support for the latest hardware and several other performance features along with the requisite bug fixes. With Windows Server 2008 R2, there are two types of installation options to choose from: Standard Installation and Server Core.

Standard: The standard installation mode is available for any release of Windows Server 2008 R2 from Standard to Datacenter. The different releases include different levels of "guest licensing rights" which varies from only one instance of a hosted Microsoft operating system license to unlimited virtual instances of Microsoft operating systems.

Server Core: Each of the Windows Server 2008 R2 releases also includes a special installation mode called "server core." The server core installation option reduces the attack surface of the operating system by reducing the number of installed binaries and limiting the roles the server can host. Most notably, Server Core uses 75% less disk space and the installations do not include Windows Explorer, so all administration will need to be done remotely or through a command-prompt. When using the Core installation mode, the "guest licensing rights" are still available, whereas with the Hyper-V Server 2008 R2 release, all guests must be licensed.

Technically speaking, the Server Core version has slightly better performance. In most cases, the Server Core version is recommended because of the decreased attack surface and increased performance. Using the Standard Installation is recommended when familiarity with the command-line is limited or the third-party software to be installed on the server is not supported on Server Core version. For instance, if System Center will be installed on the server the Standard Installation is required.

From a XenDesktop perspective any of the installation modes that include the Hyper-V role will provide the necessary interfaces. The server core role is most preferred since it provides the best performance. For more information about the different Hyper-V roles and editions, see Microsoft's Hyper-V website at http://www.microsoft.com/hyper-v-server.

Hardware Considerations

When looking for hardware that will host the virtual desktops, look for hardware that can take advantage of all the new features supported by Windows Server 2008 R2. Here are some of the hardware-based features that improve Hyper-V performance:

- Second-Level Address Translation (SLAT) enabled processors. In Intel Processors, this feature is referred to as Nested Page Tables and is available in Nehalem 55xx series processors. For AMD, this feature is referred to as Enhanced Page Tables (EPT) and is available in the Opteron 2356 Barcelona series processors. Using SLAT-enabled processors will provide the best scalability possible for the Hyper-V hosts.
- Virtual Machine Queues (VMQ) enabled network cards, especially on 10GB NICs.
- CPU Core-Parking (Processor feature of Nehalem) which allows the Hyper-V server to use the least amount of processors to meet the workload demand.

- If running Windows XP as the guest OS, be sure the processors selected support Flex Priority.

Desktop High-Availability (Clustering)

High availability is another key decision when designing the XenDesktop farm. If the user desktops need to be highly available and require Live Migration support, then Hyper-V clustering will need to be configured in the environment for those desktops. Since clustering increases overhead to the management of the virtual machines and the hypervisor, Microsoft recommends that failover clustering not be used unless high-availability of desktops is a customer requirement.

When clustering is used the following limitations will be imposed:

1. Live migration must be manually initiated and is not automatically initiated by the management system or XenDesktop.
2. Live migration of a desktop must occur between two hosts within the same cluster.
3. Live migration requires a shared volume be used to hold the virtual machine data that is to be migrated between the two hosts.
4. Live migration will only move a single desktop at a time.
5. When in a Hyper-V cluster, the supported limit of virtual guests is reduced from 384 to 64 guests. Microsoft FAQ provides this information http://www.microsoft.com/hyper-v-server/en/us/faq.aspx.
6. The maximum cluster size is 16 nodes and therefore the maximum number of guests is limited to 1024 per cluster node. Recommended load is at 15 servers (with one spare) so most clusters will support about 960 guests.

Essentially a cluster is a set of up to 16-servers that share disk and network resources. Windows Server 2008 R2 clustering comes with a "cluster validation" wizard to verify the hardware, networks, and storage systems are configured correctly. Once the cluster validation wizard completes successfully, the clustered shared volumes can be configured.

The number of clustered shared volumes (CSV) needed for the cluster depends primarily on the I/O profile of the users and the performance of the shared storage used. In most situations, two clustered shared volumes for a 16-node cluster should be sufficient. One thing to consider is the risk associated with a single host going to redirected I/O mode.

Redirected I/O mode is where all the writes from a single Hyper-V node get redirected over the network instead of written directly to the CSV. When a cluster node is in redirected I/O mode, the performance of the guests is severely impacted on both that node and the CSV owning node, since the owning node becomes a disk bottleneck. The following events can put the cluster into redirected I/O mode:

- The cluster node loses connection to the shared storage. This could be caused by any number of situations from hardware failure to storage reconfiguration.
- Backing up the CSV may cause redirected I/O for the duration it takes the back up of the volume to complete.
- Any node running Defragment or ***chkdsk*** on the CSV
- Administrators manually placing the CSV in maintenance mode

The best approach to determining the appropriate number of clustered shared volumes for the system is to conduct an I/O profile test for anticipated user traffic and use that data to calculate the amount of IOPS each volume will need to handle. With that information and the storage vendor specifications, the correct number of volumes can be calculated. The minimum number of volumes supported is one. The maximum number is 16 or one per cluster node.

CPU Considerations

Microsoft supports up to an 8:1 vCPU to physical core oversubscription ratio. This means that if the host has 4 physical cores, the number of single vCPU guests should be kept to a maximum of 32. Although no architectural limits exist, the oversubscription ratio of eight virtual CPUs to a single physical core is the current supported density. Microsoft is currently reviewing this recommendation and may increase the supported oversubscription ratio in the near future.

The only physical limitation is that a single Hyper-V host server cannot exceed 512 total vCPUs for all guests and the maximum number of guests per host is 384.

Networking

Windows Server 2008 R2 adds support for TCP chimney offload, Virtual Machine Queues (VMQ) and jumbo frames. This support improves network performance and decreases CPU utilization, thereby increasing the overall system capacity.

TCP chimney offload allows the OS to offload the processing of a TCP/IP connection to the network adapter directly. This feature allows network stack processing to be offloaded to the network adapter. See the Microsoft KB article 951037 for more information on how to configure this when the network adapter supports it. (http://support.microsoft.com/kb/951037)

VMQ-enabled network adapters allow a direct route between the virtual machine and the hypervisor by creating a unique network queue for each virtual network adapter and linking to it via the virtual machine memory. This direct route saves system resources and CPU cycles.

Jumbo frame support needs to be enabled on the core networking equipment and within the Hyper-V network adapter driver settings. Enabling jumbo frames will speed the network performance when transferring large amounts of data; however, jumbo frames do not work in all environments. Check with your network administrator before implementing jumbo frame support.

When NIC Teaming is enabled, the OS disables the TCP Chimney Offload and VMQ functionality since it cannot determine which adapter will be managing the connection. Keep in mind that if NIC Teaming is enabled, the guest has fault-tolerance for that network. If NIC Teaming is not enabled, the adapter can leverage VMQ and TCP Offload, but guest will have no network fault-tolerance for that virtual network.

When using clustering, all nodes must have the same number of network interfaces and the interfaces must be connected to the same networks. Clustering designs usually involve five networks. The standard configuration for the five networks is listed below.

- Network 1: SCVMM and Hyper-V management traffic
- Network 2: Internal cluster heartbeat
- Network 3: Public - Host/Guest external traffic.
- Network 4: Live Migration (cluster communication) and cluster metadata.
- Network 5: Dedicated for SMB/Redirected I/O traffic

If the servers are network-constrained, the minimum number of recommended networks is three. Networks 1 and 2 could be combined for management and Networks 4 and 5 can be combined for data transfer if Live Migration will not be used at all or infrequently.

Virtual Machine Manager

System Center Virtual Machine Manager (VMM) is the interface that allows XenDesktop to communicate with the Microsoft Hyper-V hosts. VMM provides a service-based interface that can be used by the XenDesktop Pool Management service to manage the virtual desktops. The VMM Deployment and VMM Operations guides can be downloaded from http://go.microsoft.com/fwlink/?LinkId=162764. Most of the high-level guidance provided in this section comes from those two guides. For more detailed information, please refer to those guides or the System Center Technet website found at http://technet.microsoft.com/en-us/scvmm/default.aspx.

Design Recommendations

Microsoft Virtual Machine Manager includes four components: VMM Server, VMM Library, VMM Database, and VMM Console. In smaller installations of less than 20 hosts these components can normally be combined on a single server. For installations of over 20 hosts the VMM server and the VMM Database should be separated.

When using a separate database server for hosting the VMM database, be sure to do the following:

- Enable remote connections on the SQL Server. More information on configuring remote connections can be found at http://go.microsoft.com/fwlink/?LinkId=127719.
- Configure the Windows Firewall to allow SQL Server Access. More information on configuring the firewall can be found at http://go.microsoft.com/fwlink/?LinkId=128365.
- Configure the SQL Server service account to run under local system or create a Service Principal Name (SPN) for the Network Service or Domain Account used for the SQL Service. More information on configuring SPN can be found at http://go.microsoft.com/fwlink/?LinkId=88057.

When designing a VMM 2008 R2 infrastructure, keep in mind the following facts:

- VMM 2008 R2 is required to manage Windows Server 2008 R2 Hyper-V hosts.
- VMM Server cannot be installed on a Server Core or a Hyper-V 2008 R2 installation.
- All VMM components are supported in a virtualized environment.
- Multiple VMM Administrative Consoles open at the same time add overhead to the management system. Keep unnecessary Administrative consoles closed.

For reference, Microsoft has verified the VMM Server can manage 400 Hyper-V unclustered hosts and 8000 virtual machines on a physical quad-core server with 8GB of RAM. The maximum clustered environment verified was 16 hosts and 1024 virtual machines on the same hardware.

Microsoft recommends when managing 1000+ virtual machines with VMM Server that the VMM Database be separated and that the SQL Server hosting the database only supports a single VMM Server. The SQL Server should have a minimum of four cores and 16 GB of RAM and use high I/O storage for the database.

Integrating with XenDesktop

Within the XenDesktop Admin console, an individual desktop group can be associated with only a single VMM Server. However, that VMM Server can service multiple desktop groups.

Since any Desktop Delivery Controller (DDC) can assume the Pool Master role, PowerShell and the VMM Administrative console must be installed on each controller in the farm. With XenDesktop 4, the DDC must be installed on Windows 2003 server so the Windows Server 2003 R2 64-bit version is recommended. The VMM Administrative console cannot be installed on a 32-bit version of the Windows Server 2003, but it can be installed on a 32-bit version of Windows Server 2003 R2.

Finally, on the VMM Server, disable VM auto-start since it will interfere with XenDesktop's management of the machine state.

Configuring for Disaster Recovery

Microsoft's VMM server does not support a farm concept for high availability; instead, each VMM server is a self-sufficient entity. The only fault-tolerant solution available for VMM is to use Microsoft clustering.

VMM Server: Use a Microsoft Hyper-V failover cluster and create a virtual server that is configured for high-availability. If the VMM Server is virtualized, do not manually migrate it between hosts or configure it for Performance Resource Optimization (PRO) which could inadvertently cause a migration.

VMM Database: Use a Microsoft failover cluster to host SQL Server and place the VMM Database on the clustered SQL Server. Using SQL Express as a local database is not recommended in environments with more than five Hyper-V hosts.

VMM Library: VMM supports adding highly available library shares on a failover cluster created in Windows Server 2008 Enterprise Edition or Windows Server 2008 Datacenter Edition. Use a Microsoft Windows Server 2008 failover cluster because VMM cannot recognize a failover cluster on Windows Server 2003.

However, when the VMM Server is configured as above it only protects against hardware failure, not against software failure. If the VMM Server service hangs or stops responding to requests, the server becomes unavailable. Protecting against this type of failure will require custom detection scripts and complex failover routines.

When the VMM Server service becomes unavailable, the XenDesktop Pool Management service takes those desktops offline. Limiting the number of virtual machines managed by a single VMM server reduces the exposure in case of a software failure.

Storage Selection

The storage selected to host the VDI environment can have a huge impact on performance of the virtual desktops. Understanding the characteristics and behavior of storage with XenDesktop and Hyper-V is essential to successfully designing a storage solution.

Storage Types

When planning the Hyper-V deployment, keep in mind that Hyper-V supports only direct attached storage or block-level SAN storage such as Fiber Channel or iSCSI. Any storage selected should support approximately 15 IOPS (I/O Operations Per Second) per desktop on the host. In other words, if each host will support 50 desktops, then the storage should provide 750 IOPS (at 90% writes) per host.

Virtual desktop workloads are extremely write-intensive, with some studies showing up to 90% write activity. In a write-intensive environment, RAID 1 or RAID 10 configurations incur a 2x write overhead where RAID 5 configurations incur a 4x write overhead. In most cases using RAID 1 or

RAID 10 storage configurations will provide better performance for the virtual desktop users. Table 1 summarizes the write-penalties for the various workloads.

RAID Level	Write Penalty
RAID 0	1
RAID 1 or RAID 10	2
RAID 5 (3 data 1 parity)	4
RAID 5 (4 data 1 parity | 3 data 2 parity)	5
RAID 5 (5 data 1 parity | 4 data 2 parity)	6

Table 1: IOPS Write Penalty for RAID Configurations

The formula below can be used to calculate the functional IOPS available for a VDI workload.

Functional IOPS = ((RAW Storage IOPS * .9)/ Write Penalty) + (RAW Storage IOPS * .1)

For example, eight 72GB 15K SCSI3 drives in a RAID 10 storage array would have a total RAW IOPS of 1200 and Functional IOPS for VDI of 660.

Functional IOPS = ((1200 * .9) / 2) + (1200 * .1) = 540 + 120 = 660.

The number of desktops supported by this storage array can be calculated by dividing the Functional IOPS by 15 to get 44.

Saved State Files

With Microsoft Hyper-V, the first time a virtual machine is started the hypervisor creates a .BIN file that is equal in size to the RAM configured for the virtual machine. This file is used to store the contents of the RAM when the virtual machine is moved into a saved state. The file is created automatically in the same location where the virtual machine resides and cannot be disabled. When calculating the necessary disk space on a SAN to support virtual machines, include the amount of virtual machine RAM in the calculations.

Operating System Delivery

This section includes guidance around delivering the operating system to the virtual machines. The areas discussed include operating system selection, virtual hard disks, and networking.

Operating System Selection

Internal testing from both Citrix and Microsoft currently indicates that on Microsoft Hyper-V, Windows 7 provides a better user experience than Windows XP and generally provides better scalability when configured with the same amount of RAM.

The Microsoft team attributes this improved scalability to Windows 7 being a virtualization-aware operating system. Windows 7 includes several features which improve its performance in a virtualized environment.

- Windows 7 includes the Hyper-V Host Integration Services as part of the base operating system
- Windows 7 notifies the hypervisor when it is idle so the hypervisor does not schedule guest operations
- Windows 7 includes optimized device drivers for network and disk
- Windows 7 provides improved storage and optimized page file management

Another decision point when selecting the guest operating system is whether the virtual desktop's operating system supports all the customer applications. If some applications only execute on Windows XP or Vista or only support 32-bit operating systems, those considerations will drive a particular operating system selection. Consider using Microsoft App-V or Citrix XenApp to deliver applications not natively supported on Windows 7.

Virtual Hard Disk Considerations

In the Citrix and Hyper-V virtual environment, virtual hard disks (VHD) are used by both Provisioning Services and Microsoft Hyper-V. When working with virtual hard disks, remember to always avoid dynamic VHD files and manually set the partition offset with ***diskpart*** utility. See the Microsoft KB article 300415 (http://support.microsoft.com/kb/300415) for ***diskpart*** usage.

Avoid Dynamic VHDs

Dynamic VHDs include a single byte at the end of the file that causes the file to be out of alignment with the disk subsystem. When a file is out of alignment, it causes the disk subsystem to perform extra I/O operations for each file change and degrades performance considerably. Using a fixed disk VHD will not include the extra byte at the end of the file.

Dynamic VHDs also have an expansion algorithm (in 2MB blocks) that generates significant overhead on the storage device when the drive expands. As the dynamic drive expands the

allocation table is updated and the drive's header and footer sections are rewritten (causing a read & write) for each of the file extension operations. In other words, if a 10 MB file is written to the dynamic VHD, it will cause 30 separate I/O operations in overhead just for the management of the dynamic VHD file.

When creating the write-cache drives for Provisioning Services, use only fixed-size VHDs. Using fixed-sized VHDs for the write-cache drives avoids the bottleneck of file expansion and improves the SAN alignment prospects.

Manually Configure VHD Partition Offset

Another guideline to improve system performance is to manually set the partition offset to 1024 KB using the ***diskpart*** utility during installation. Setting the offset to 1024 KB will improve the alignment with the disk storage subsystem. When the VHD file has the correct offset, the storage system will use less IOPS when reading or writing to the file since fewer disks will be accessed.

The Windows XP setup and ***diskpart*** utilities create the boot partition with an offset that is misaligned with the disk subsystems. The ***diskpart*** utility included with Windows Server 2003 SP1 or Windows 7 has been corrected to create the boot partition at a more disk-friendly offset, but still it may (depending on the size of the partition) create it with a misaligned offset. The recommended approach is to always manually create the partition offset with ***diskpart*** on all VHD files before formatting them. For more information on calculating the correct offset see the Microsoft KB article 929491 (http://support.microsoft.com/kb/929491).

When converting an image to a VHD for use with Provisioning Services, use XenConvert 2.1 or later which will allow configuration of the VHD partition offset. XenConvert 2.1 can be downloaded at http://www.citrix.com/English/ps2/products/subfeature.asp?contentID=1860675. This version of XenConvert allows the administrator to specify the exact offset for the vDisk partition. Previous releases set the offset to 252KB which would be out of alignment with the disk storage system. More information on using XenConvert 2.1 can be found in the Appendix.

Networking Considerations

Microsoft Hyper-V has two types of network adapters. The first is referred to as the "Legacy Network Adapter" in the Hyper-V management console and as the "Emulated Network Adapter" in the VMM Administrative console. The other adapter is referred to as the "Synthetic Network Adapter" in both consoles.

The legacy network adapter is tied directly to the BIOS of the virtual machine. Using the legacy adapter increases processor overhead because device access requires context switching for communication. The legacy network adapter is required for supporting any Pre-boot Execution Environment (PXE) such as that used with Provisioning Services. Contrary to popular belief, the legacy network is not limited in speed to 100MB, but it can run at speeds higher than 100MB if supported by the host's physical network interface.

The synthetic network adapter is loaded by the Host Integration Services after the operating system loads inside the virtual machine. As such, the synthetic network adapter is not available for any PXE operations. Since the synthetic network adapter is integrated directly with the virtual machine, it can leverage the high-speed VMBus for communication and reduce processor overhead by avoiding the context switches that the legacy network adapter requires.

Single Network Adapter Configuration

If Provisioning Services or other third-party PXE imaging delivery applications will not be used in the environment, the legacy network adapter is not necessary. Best performance will be achieved by using a single synthetic adapter.

Conversely, if the processors are not taxed and can easily handle the additional context switches, the legacy network adapter could be the sole network adapter for virtual machine. The network throughput would be the same from the perspective of the virtual machine. The only impact might be on the number of the guest virtual machines supported by a single physical Hyper-V host.

A single network adapter is recommended for simplicity if network performance or fault-tolerance is not a priority for the user environment. If NIC Teaming will not be configured for the adapters servicing the virtual network or if network performance is a key requirement, the dual network adapter approach is recommended

Dual Network Adapter Configuration

With Hyper-V the legacy network card (NIC) is required for supporting PXE with Provisioning Services. After the virtual machine boots, the synthetic NIC has precedence over the legacy network card since the driver sets the route metric for the synthetic NIC to be lower than the legacy NIC.

If using Provisioning Services to deliver the operating system, be sure to run ***bindcfg*** from the Provisioning Services installation folder to verify the legacy network adapter is bound to the Provisioning Services device driver before creating a vDisk image. If the Provisioning Services device driver binds to the wrong network adapter, the image will not be able to complete the boot process. More information on using ***bindcfg*** can be found in the Appendix.

If using Provisioning Services to stream the operating system, the best performance is achieved by creating two network cards for each virtual machine. The legacy network card will be used to support the PXE booting and all PVS traffic and the synthetic network card will be used for all other network traffic once the operating system has started. If both network cards are enabled in the operating system and on the same subnet, the synthetic card should have preference for all non-PVS traffic. The PVS traffic will always traverse the legacy network card because it is bound to that card. In some situations the legacy network adapter might also be used to transmit data since Windows Networking uses multiple factors to determine the best route for a packet.

Decision Summary Table

Table 2 summarizes the design decision points and recommendations discussed in the document. This table can be used as a handy reference when architecting the XenDesktop on Microsoft Hyper-V environment.

Decision Point	Options	Recommended
Windows Server 2008 Installation Mode	Standard, Server Core, Hyper-V Server 2008 R2	Server Core
Hardware Feature Set	SLAT, VMQ, PP, Flex Priority	SLAT, VMQ
Desktop High-Availability	Hyper-V Failover Clusters or Normal Hyper-V Servers	Normal Hyper-V servers unless HA is required.
Clustered Shared Volumes	1-16	2-4 depending on user I/O Profile
Clustering Network Design	Internal, External, Management, Migration, SMB (1) – (5)	(3) External, Mgmt/Internal, Migration/SMB
Network Configuration	Jumbo Frames, VMQs, TCP Offload, NIC-Teaming	VMQs, TCP Offload
CPU Oversubscription	Calculate the number of vCPUs for all guests (maximum 384) and divide it by the number of physical cores	8:1 with a maximum of 512 vCPUs
VMM Server Configuration	Separate VMM Database or single-server installation	Separate VMM Database from VMM Server
VMM Server Virtualized	Virtual or Physical	Virtual (can be clustered for HA)

Decision Point	Options	Recommended
Storage calculations	IOPS per desktop Saved State files	Minimum IOPS 15 per virtual machine. Include VM RAM in total storage requirements
SAN Storage RAID Level	RAID1, RAID5, RAID10	RAID1 or RAID10
Desktop Operating System	Windows XP, Windows Vista, or Windows 7	Windows 7
VHD Type	Fixed or Dynamic	Fixed
VHD Partition Offset	Automatic or Manual	Manually configure partition to match storage blocks
Guest Network Adapters	Single or Dual	Single if possible Dual for performance

Table 2: Hyper-V Design Decision Summary

Deployment Considerations

This section covers some deployment considerations for large environments. The first section includes guidance for XenDesktop DDC settings. The next two sections provide guidance and scripts for creating and assigning the Provisioning Services write-cache drives in bulk.

Desktop Delivery Controller Settings

In large desktop farms, tuning the Desktop Delivery Controller (DDC) performance and creating dedicated roles is recommended. To change the DDC roles see Citrix Knowledgebase Article CTX117477 http://support.citrix.com/article/ctx117477 . See the Appendix for more guidance on setting roles.

Current recommendations for architecting XenDesktop includes dedicating a master for brokering and pool management then using the remaining member servers for XML requests. In other words, when configuring Web Interface, place the Farm Master last in the list of servers for Web Interface to contact for the farm.

If any DDC will need to handle more than 3900 connections, a registry change will need to be made on the DDC because by default Windows Server 2003 is limited to 3977 TCP connections. (Server 2008 defaults to 16,384) For more information see http://support.microsoft.com/kb/196271.

XenDesktop Creation

The XenDesktop Setup Wizard does an excellent job of interfacing with Provisioning Services, the VMM Server, and the XenDesktop controller to build all the desktops automatically. However, it creates desktops at the rate of two per minute for clustered Hyper-V servers and three per minute for unclustered Hyper-V servers.

One approach is to run multiple instances of the XenDesktop Setup Wizard on the Provisioning Server. The XenDesktop Setup Wizard must be run from the Provisioning Server itself, since a remote API for the Provisioning Server commands is not yet available. Running multiple instances of the XenDesktop Setup Wizard on the same Provisioning Server or multiple Provisioning Servers in the same farm will work well as long as the VMM Servers targeted are different.

Running multiple instances of the XenDesktop Setup Wizard against the same VMM Server is not recommended because the Setup Wizard does not use the VMM APIs for intelligent placement. Instead the Setup Wizard takes the number of desktops to be created and equally divides the machines across all the hosts registered with the VMM Server.

A second approach is to replace the XenDesktop Setup Wizard functionality with PowerShell scripts that access the VMM Server, Provisioning Server, and XenDesktop APIs. All the components support an API for configuration; however, at this point the replacement scripts are not available.

When the Hyper-V servers have unequal capacity such as in cases where the servers are already partially loaded or have different capacities, using a VMM Staging server will allow the XenDesktop Setup Wizard to be used. The staging server functions as a temporary VMM management server for the Hyper-V hosts. To use a VMM staging server, follow these high-level steps:

1. Build a VMM Server with a local SQL Express database.
2. Temporarily register the Hyper-V hosts with the Staging VMM Server.
3. Run the XenDesktop Setup Wizard to create the desktops.
4. After the Setup Wizard completes, re-register the Hyper-V hosts with the production VMM Server
5. Modify the XenDesktop Group properties to use the production VMM Servers.

Use the staging server to avoid writing scripts for virtual machine creation and placement.

Assigning write-cache drives in bulk

Another side-effect of using the XenDesktop Setup Wizard is that any drives attached to the virtual machine used as a template are discarded from the creation, so all write-cache drives need to be added after the desktops are created. The VMM Server can be used to create the write-cache drive and attach it to the virtual machine. The Appendix contains a PowerShell script that can be run from the VMM server to attach a VHD from the SCVMM Library to all the machines that match a naming pattern provided. This script will take about 30-90 seconds for each virtual machine or possibly longer depending on the storage and network speeds. Furthermore, it is run sequentially by the VMM Server.

Another, and probably faster way, to do this is to write a batch file that copies a template VHD (partitioned with a 1024 KB offset and formatted appropriately) for each virtual machine then run a PowerShell script from the VMM Server to just attach the VHDs after they are created. With this method, the file copies which take the longest amount of time can be run in parallel from each of the hosts and if correctly configured can take about 20 seconds for the file copy.

Keep in mind to avoid the "new hardware found" warning message the write-cache drives should be a copy of the write-cache VHD attached to the PVS vDisk image, that way the disk signatures will match and Windows will not prompt the user for a restart.

Conclusion

This Design Guide was created to provide specific guidance when designing a XenDesktop deployment on Microsoft Hyper-V. This guidance is to be used in conjunction with other XenDesktop architecture guides. For more information on designing a XenDesktop deployment, see the following whitepapers and URLs.

XD Design Handbook: http://support.citrix.com/article/CTX120760

Delivering 5000 Desktops with XD4: http://support.citrix.com/article/CTX123684

XenDesktop Modular Reference Architecture: http://support.citrix.com/article/CTX124087

XenDesktop Whitepapers: http://support.citrix.com/product/xd/v4.0/whitepaper/

XenDesktop on Microsoft Website: http://community.citrix.com/p/edit/xendesktop-on-microsoft

Appendix

The Appendix is designed to provide additional technical instructions beyond what was appropriate for the main document content.

Configuring Partition Offset Manually with XenConvert 2.1

The XenConvert 2.1 documentation describes how to manually change the partition offset, but for simplicity it is described here as well. The offset of the partition can be set by modifying the XenConvert.INI file found in the same folder as the XenConvert executable. Add the following section and value to the file to set the offset to 1024 KB.

```
[parameters]
PartitionOffsetBase=1048576
```

With this INI file parameter you can match the partition offset to the exact size necessary for any local or SAN storage system. Use this procedure for all VHDs that will reside on the SAN storage system.

Using BindCfg to Set the Network Card

If using multiple network adapters in the host, verify the correct network adapter is bound to the Provisioning Services stream service. To verify the Provisioning Services network card binding, execute the \Program Files\Citrix\Provisioning Services\BindCfg.exe program on the XenDesktop. Verify the Hyper-V Legacy (Emulated) Network Adapter is selected as shown in the screenshot below. The Microsoft Virtual Machine Bus Network Adapter is the Synthetic network adapter that does not support PXE booting.

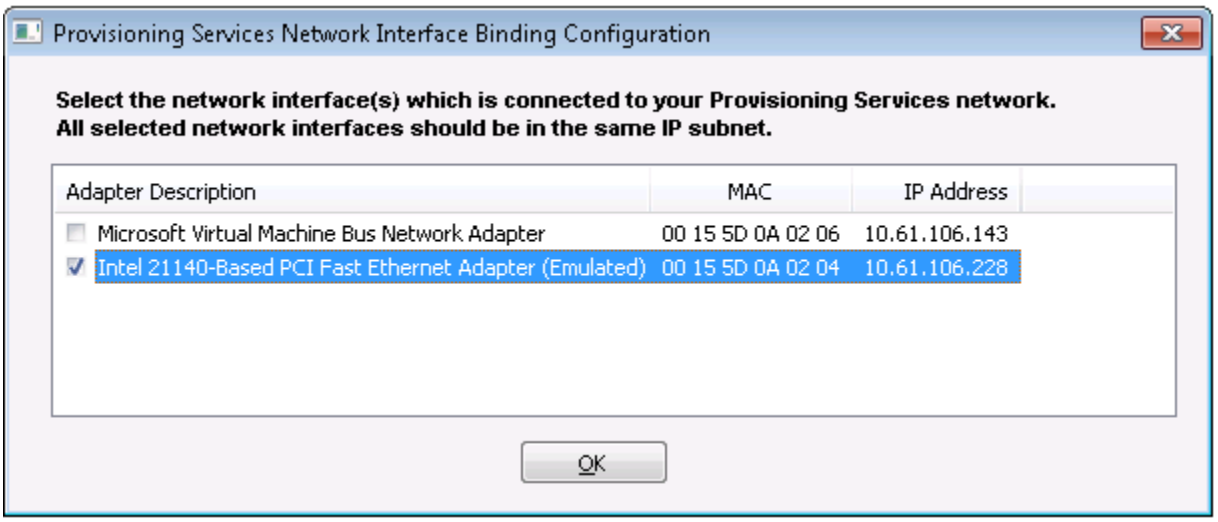

PowerShell Script to Copy and Attach the VHD from VMM

The following PowerShell script can leverage VMM Server's API to both copy a VHD from the VMM library to the host and then attach it to an existing virtual machine.

Sample Usage: PS c:\scripts> .\copyvdisk.ps1 "\\SCVMM\MSSCVMMLibrary\VHDs\writecachebase.vhd" "C:\VMs" XDesktop "" ""

Requirements:

1. Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded
2. The base VHD must be included in the VMM Server's library. May require a refresh on the SCVMM console. The VHD should also be partitioned and formatted with NTFS.
3. The first parameter is the path to the VMM library and must be a UNC path.
4. The second parameter is the path to the location of the VM's on the Hyper-V server itself and it must be relative to the Hyper-V host and not a UNC path.
5. The third parameter is the naming pattern to match with an assumed wild-card at the end of it.
6. The fourth and fifth parameters are prepend and append strings to the write-cache VHD filename which without any modifiers will be equal to the virtual machine name.

COPYVDISK.PS1 contents:

```
if ($args -eq $null -or $args.Count -lt 3)
{
	write-output "Usage: copyvdisk UNC_fileName_of_vhd Path_to_cluster_storage VMNameMatches prepend_filename
append_filename"
	exit 1
}
$VdiskPath = $args[0]
$ClusterStoragePath = $args[1]
$VMNameMatches = $args[2]
$PrependString = $args[3]
$AppendString = $args[4]

$VMMServer = Get-VMMServer -Computername "localhost"
$BaseVdisk = Get-VirtualHardDisk -VMMServer $VMMServer | where { $_.Location -eq "$VdiskPath" }

if ($BaseVdisk -eq $null)
{
	write-output "Unable to find vdisk: $VdiskPath"
	exit 1
}

$VMs = Get-VM | where { $_.Name -match "$VMNameMatches" }
if ($VMs -eq $null)
{
```

```
        write-output "No VMs matches the pattern: $VMNameMatches"
        exit 1
}
else
{
        $matchedString = "{0} vms match the pattern: {1}" -f $VMs.Count, $VMNameMatches
        write-output $matchedString
}

foreach ($vm in $VMS)
{
        $current_disks = get-VirtualDiskDrive -VM $VM
        if ($current_disks -eq $null -or $current_disks.count -eq 0)
        {
        $filename = "{0}{1}{2}.vhd" -f "$PrependString", $VM.name, "$AppendString"
        $cloningMessage = "Attaching {0} to VM {1}" -f $filename, $VM.Name
        write-output $cloningMessage

        $newvhd = New-VirtualDiskDrive -VM $VM -VirtualHardDisk $BaseVdisk -Path "$ClusterStoragePath\$VM" -Filename
"$filename" -IDE -Bus 0 -LUN 0
        }
        else
        {
                $diskattachedmessage = "{0} {1}" -f $VM.Name, "has disk already attached"
                write-output $diskattachedmessage
        }
}
```

Configuring a Dedicated Farm Master

For the most part, following the instructions outlined in the Citrix Knowledgebase Article CTX117477 (http://support.citrix.com/article/ctx117477) can be followed to set the registry keys appropriately for the farm servers. The only steps not covered in the article is configuring web interface to only use the specific servers for XML requests.

To edit the servers used for the XML requests, go to the Web Interface Management console (on all Web Interface servers servicing the farm) and complete the following steps:

1. Select XenApp Web Site or XenApp Services Site that is being used.
2. Right-click the site and choose **Server Farms** from the context menu.
3. Select the Farm and click the **Edit…** button.
4. The dialog box shown in the screenshot below will then be available. In the screenshot below, *xdddc2-xdddc4* function as the XML request servers and xdmaster is the farm master. Notice that *xdmaster* is still listed, but will only be used after all other farm servers are unavailable.

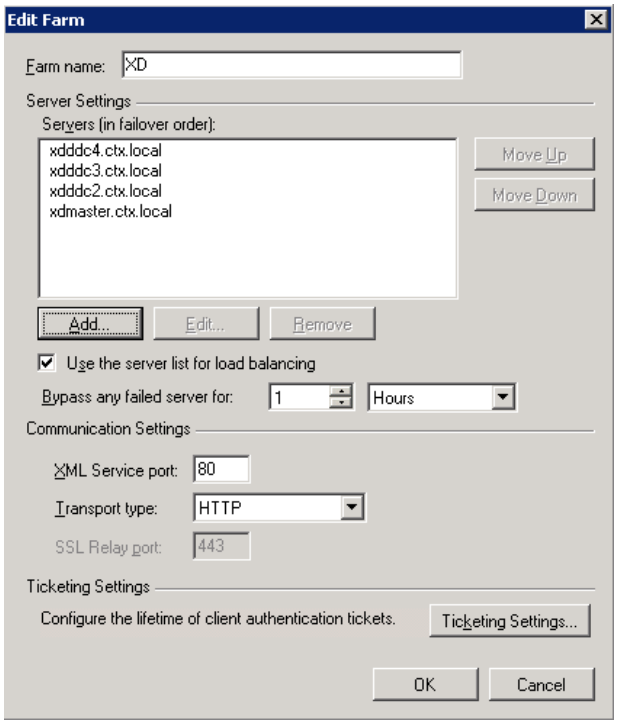

Citrix Consulting - Americas
Citrix Systems, Inc.
851 W Cypress Creek Road
Fort Lauderdale, FL 33309 USA
Phone: 1-800-393-1888 Option 55
Phone: 1-954-229-6170

Citrix Consulting - Pacific
Citrix Systems Singapore Pte Ltd

8, Temasek Boulevard
#34-02, Suntec Tower 3
Singapore, 038988
Singapore

Citrix Consulting - Europe
Citrix Systems UK, Ltd.
Chalfont House
Chalfont Park, Gerrards Cross
Bucks SL9 0BG
Phone: +44 1753-276-200

Citrix Systems Japan KK

Kasumigaseki Common Gate West Tower 24F
3-2-1, Kasumigaseki, Chiyoda-ku
Tokyo, 100-0013 Japan

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. Its Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2008 was $1.6 billion.

©2010 Citrix Systems, Inc. All rights reserved. Citrix®, Access Gateway™, Branch Repeater™, Citrix Repeater™, HDX™, XenServer™, XenApp™, XenDesktop™ and Citrix Delivery Center™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.