

Cisco Unified Computing System with Microsoft Hyper-V Recommended Practices

White Paper
January 2014



Preface

As both Cisco Unified Computing System™ (Cisco UCS®) and Microsoft Hyper-V continue to gain market share, more people are requesting information about recommended practices around combining these technologies. The purpose of this paper is to provide a series of recommendations based on the experience of individuals who have been working with Cisco UCS and/or Microsoft Hyper-V for many years.

Table of Contents

Introduction	2
Hyper-V Hosts	2
Types of Installations	2
Management	3
SAN Boot	3
Host Networking	4
Configuring NICs with VLANs	4
NIC Binding Order	5
Microsoft NIC Teaming versus Fabric Failover	6
Quality of Service	7
Management IP Addresses	8
Host Storage	8
Prepare for Virtual Machine Storage Access	9
Host Clustering	10
Storage	10
NICs	10
Network Prioritization	11
Host Memory	12
Host Memory Calculation	12
Paging File	13
Virtual Machines	14
Virtual Machine Generation	14
Virtual Machine Networking	15
NICS	15
Cisco Nexus 1000V	15
SR-IOV versus VM-FEX	16
Virtual Machine Storage	16
Virtual Fibre Channel	17
Virtual Machine Clustering	17
Shared Storage	17
Summary	18
About the Author	20
Acknowledgements	20

Introduction

Both Cisco Unified Computing System (Cisco UCS) and Microsoft Hyper-V continue to grow in market share. It is only natural that more and more organizations are asking for recommendations on how to deploy Hyper-V on Cisco UCS. There are many things that are common recommendations no matter what server platform is being used. But Cisco UCS brings some unique technologies into the environment, presenting different ways to architect a solution.

For example, Microsoft offers network interface card (NIC) teaming capabilities as part of the operating system, and one feature of NIC teaming is aggregated bandwidth. Cisco offers Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX), a Single Root I/O Virtualization (SR-IOV) solution that can be used to aggregate bandwidth. Microsoft's NIC teaming solution also provides failover capabilities. Cisco offers virtual port channels and port channel groups that provide failover capabilities. Which technology should be used? Is one more appropriate than the other? Is one technically superior to the other?

The goal of this document is not to try to position one technology over the other to say one is technologically superior, but rather to explain the pros and cons of using one or the other in various solutions. The goal is to give you enough information to make the decision for your particular environment.

Please note that this document contains *recommendations*, not hard and fast rules. Specifics of your environment are unknown here. Some of these recommendations may make sense. If so, implement them. Some of them may not bring any value to your environment. In that case, ignore them.

Source of these recommended practices come from many sources and experts. Roger Osborne is a support engineer for Microsoft. He assembled a list of general practices for Hyper-V and published it in a [blog](#). Since it was a nicely ordered list, it is recommended you look at that list for general Hyper-V settings. Other information comes from personal experience with Hyper-V as well as other experts at Cisco and Microsoft. This document will focus on specific capabilities of Cisco UCS, but recommend you use it in conjunction with Roger's list, which expands some of the recommendations in this document.

Additionally, Cisco has published [Cisco UCS Manager Configuration Common Practices and Quick-Start Guide](#). This guide provides an overview of Cisco UCS essentials and best practices. In general, the guidelines provided in that document should be followed, but there may be some variations to consider when implementing in a Microsoft environment. A goal of this document is to explain these variations.

Lastly, there are some general Hyper-V or Windows Server configuration recommendations that are not covered in any of the resources we've mentioned so far. These may not be unique to running Hyper-V on Cisco UCS, but they are recommended practices for any vendor's platform. The combination of the document at Roger Osborne's site and the Cisco UCS Manager guide with this document should provide you with most of the recommendations for successfully deploying Hyper-V on Cisco UCS.

Hyper-V Hosts

Types of Installations

You can use Hyper-V in a variety of different installations. No matter how it is installed, it is the same Hyper-V. It comes as a role in Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and even as an option on Microsoft's 64-bit Pro and Enterprise Editions of Windows 8 and 8.1. This document focuses on running a server version on Cisco UCS; it does not address a client implementation, though many of the recommendations would be the same.

Obtaining Hyper-V with the purchase of the Windows Server operating system grants specific virtualization rights for virtualizing instances of Windows Server. Hyper-V Server 2008, Hyper-V Server 2008 R2, Hyper-V Server 2012, and Hyper-V Server 2012 R2 come as no-cost downloads from [Microsoft](#). This installation does not come with any virtualization rights, so licenses for virtual machines running on this installation must be purchased separately.

Even though Hyper-V is the same across all platforms, there will be some slight variations in capabilities when running on the client operating systems because they do not support server features, such as failover clustering of the host (though clustering still works fine in the VMs running on the client version of Hyper-V). As with the no-cost Hyper-V Server downloads, no server virtualization rights are included with the client license, so any operating systems installed on the client will need to be licensed separately. Licensing is not covered in this document. Contact a trained Microsoft licensing specialist for questions regarding virtualization rights granted with different operating system licenses.

Management

Recommendation: Manage from the “top down.” Consider Hyper-V Manager as the lowest level, Failover Cluster Manager as the next level, and System Center Virtual Machine Manager (SCVMM) as the topmost level. Attempt management functions first from SCVMM, then from Failover Cluster Manager, and lastly Hyper-V Manager.

Hyper-V comes with fairly robust management tools and capabilities without any additional software. The Hyper-V Manager console allows access to all virtual machines’ settings across one or more hosts. When Hyper-V hosts are built into a Microsoft Failover Cluster environment, some of those management functions must be performed from the Failover Cluster Manager console, but they operate on the same values you would find in the Hyper-V Manager console. For an additional cost, Microsoft System Center Virtual Machine Manager (SCVMM) is the most robust management environment that allows management of hosts, VMs, networks, storage, and other facets of a virtualized environment.

In installations with multiple management consoles, it is highly recommended to manage from the top down. If you are just running Hyper-V, all management will be done through the Hyper-V Manager console (or associated PowerShell cmdlets). When you cluster Hyper-V hosts, management should first be attempted through the Failover Cluster Manager console. If required, the Failover Cluster manager console will redirect you to the Hyper-V Manager console, but there are few things that you cannot do in the Failover Cluster Manager console. If managing the environment with SCVMM, you should do all management through the SCVMM Manager console.

The reason for this hierarchy is that a higher-level management tool is reaching down to the lower levels to accomplish the end goal. As a result, the higher-level management will know what is being updated and will keep its management database up to date. If a Hyper-V host that is managed by SCVMM is altered through the Hyper-V Manager console, it does not presume to know anything about upper levels of management, so any change made may not be recognized by the higher-level managers until a refresh is done on the higher-level manager. Generally, this will not cause a major problem, but there have been times when it has taken some time to get everything back in synch. You avoid this waste of time when the system is managed from the top down.

Most of this document focuses on managing directly to Hyper-V. However, as we will note, there are some places where recommendations are made that would be slightly different if SCVMM is used for management.

SAN Boot

Recommendation: As the industry leader in Fibre Channel over Ethernet (FCoE) and converged network technologies, Cisco recommends using FCoE for SAN boot whenever it is available. Otherwise, use a Fibre Channel (FC) boot. At this time, it is not recommended to use iSCSI boot; it works, but it is not as robust or stable as FCoE or FC boot.

As the earlier referenced Cisco UCS Manager Configuration Common Practices and Quick-Start Guide explains, “Cisco UCS provides the infrastructure to support the highest levels of service availability through a stateless server environment. Stateless servers are logical servers (OS and application images) that can run on any physical server or blade in a way that encapsulates traditional hardware identifiers ... as integral parts of the Service Profile identity.” Key to this flexibility is operating system boot from a SAN logical unit number (LUN) to allow replacement of physical server hardware without any need to reconfigure the OS and application image.

Windows Server has long been able to boot from SAN. What has changed over the years are the types of SANs available. Of the variety of storage protocols supported by Windows Server, there are three that are available for SAN booting: FCoE, FC, and iSCSI.



Cisco provides the FCoE and FC boot capabilities in its adapters, providing robust solutions for either configuration Cisco does not have iSCSI boot capabilities in its adapters, so it needs to rely on Microsoft's iSCSI initiator to perform a SAN boot to an iSCSI SAN. It can be done, but we have encountered various issues with iSCSI boot, so iSCSI SAN is not a recommended solution for boot.

Windows uses LUN 0 from which to boot. Therefore, use LUN 0 for the boot LUN in the boot profile of the Service Profile.

When initially installing the operating system, configure just a single path to the LUN making use of fabric zoning and LUN masking. Until the operating system is installed and multipath I/O (MPIO) software is installed on the operating system to handle the multiple paths, data corruption on the LUN or other problems can occur. Therefore, single path during initial installation is a requirement.

Microsoft has a utility called sysprep that is used in many Microsoft-centric data centers. This utility allows for an initial installation to be completed, tailored to contain specific software settings and applications, and then prepared for redeployment. For example, Microsoft's MPIO software can be installed in the system before sysprep is run. A copy of this image can be used to very quickly deploy another instance of the operating system instead of going through the full Windows Server installation process. When an image is created with Microsoft's MPIO software already installed and configured, it is possible to configure multiple paths to LUN and boot.

Most storage vendors have a way to clone LUNs for reuse. Work with the storage vendor to set up a procedure to build a sysprep "gold image" on a LUN that can be cloned for rapid deployment.

Host Networking

Configuring NICs with VLANs

When you configure NICs on Hyper-V, there are multiple places that VLAN tags can be specified, depending on the use of the NIC being created.

Host-only NIC: When creating a NIC for a network that will be used only by the host, use Cisco UCS Manager to specify the VLAN as a Native VLAN. Cisco UCS will ensure that the packets are properly tagged.

NIC shared between host and VMs: If you are creating a NIC for a network that will be used by both the host and VMs, do not specify the VLAN as a Native VLAN within Cisco UCS Manager. Instead, when you create the Hyper-V virtual switch on this NIC, create the virtual switch to *Allow management operating system to share this network adapter* and *Enable virtual LAN identification for management operating system*. Then specify the appropriate VLAN tag, as shown in the following image.

Virtual Switch Properties

Name:
Shared_NIC

Notes:

Connection type
What do you want to connect this virtual switch to?

☒ External network:
Cisco VIC Ethernet Interface #8

☒ Allow management operating system to share this network adapter

☐ Enable single-root I/O virtualization (SR-IOV)

☐ Internal network

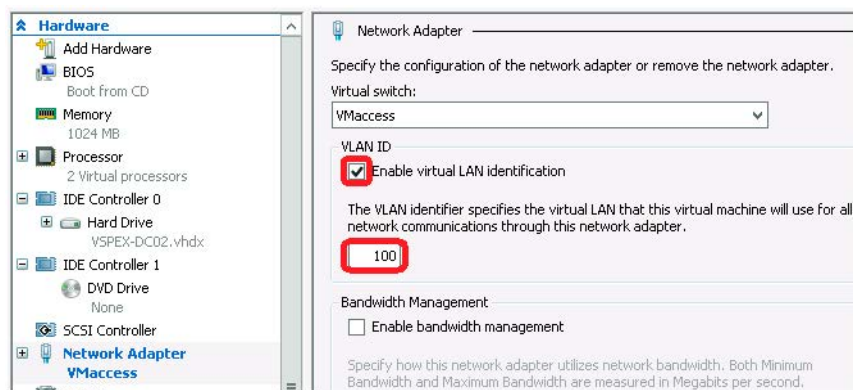
☐ Private network

VLAN ID
☒ Enable virtual LAN identification for management operating system

The VLAN identifier specifies the virtual LAN that the management operating system will use for all network communications through this network adapter. This setting does not affect virtual machine networking.

10

VM-only NIC: If you are creating a NIC for a network that will be used only by virtual machines, do not specify the VLAN as a Native VLAN within Cisco UCS Manager. When you create the Hyper-V virtual switch on this NIC, do not specify to *Allow management operating system to share this network adapter*. This will ensure that you cannot *Enable virtual LAN identification for management operating system*. The VLAN tag has to be specified within the individual VM settings; there is no global setting. The VLAN tag **must** be set here if the VLAN is to be used by the VM, as shown in the following image.



This manner of setting of the VLAN tag for the individual virtual machines is the same whether the virtual switch is configured as shared with the management operating system or not.

When you use SCVMM, it will automatically take care of the assignment of VLAN tags within the virtual machines settings.

NIC Binding Order

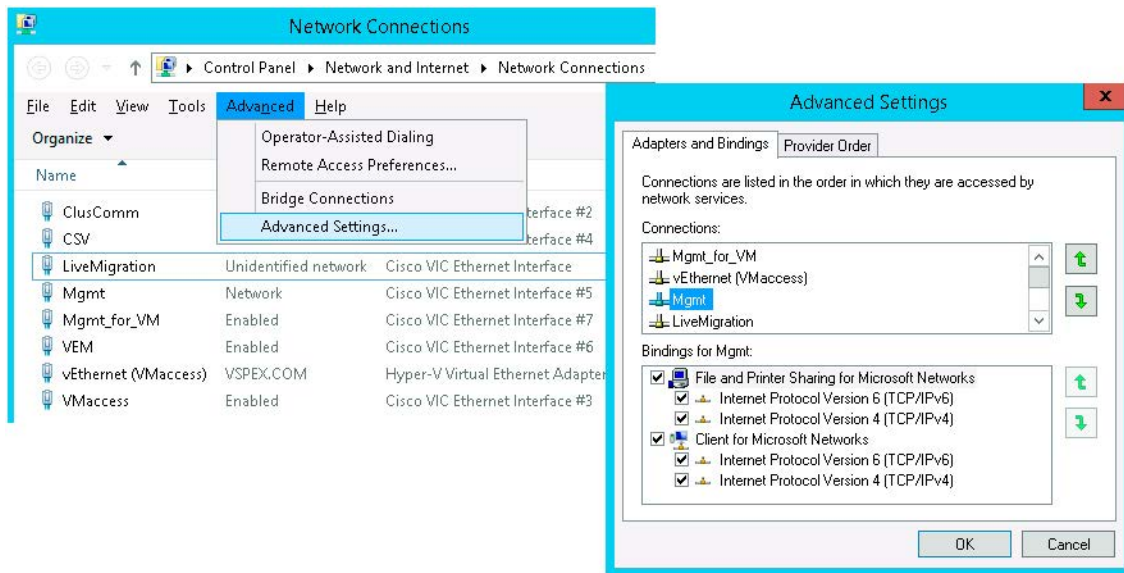
Recommendation: Ensure the primary NIC, generally the NIC on the domain network, is set first in the binding order.

NIC binding order is the order in which NICs are utilized to communicate for things like authentication and authorization. If NICs are not in an optimal order, the system will still run, but it might take longer to log in or gain authorization to a resource because a NIC that is not on the proper network is tried first and a timeout needs to occur before the next NIC in the binding order is tried.

NIC binding order is not something that is unique to a Cisco UCS installation. Deployment of Hyper-V on any vendor's platform will present the same issues. Hypervisors in production environments are always deployed with multiple NICs. As Hyper-V discovers and configures the NICs, they may end up being configured in a binding order that is suboptimal.

Once NICs are configured, it is recommended to move the primary NIC to the top of the binding order. The primary NIC is considered the NIC over used for the majority of requests for authentication and authorization. In most environments, this would be defined as the management NIC that communicates to the Active Directory domain controller.

Setting the binding order is accomplished through the Control Panel > Network and Internet > Network Connections. If the menu line is not showing at the top of the window, press the ALT key to expose the menu. Select the *Advanced* menu item, and then select *Advanced Settings...* In the Advanced Settings window, select your primary network under Connections and use the arrows on the right-hand side to move it to the top of the list. Click *OK* to accept the changes. The following images illustrate these settings.



Microsoft NIC Teaming versus Fabric Failover

Recommendation: Deploy Cisco UCS fabric failover. No additional work is required; this is part of the configuration of Cisco UCS and Cisco Nexus®. Network managers do not need access to Windows Server utilities to define NICs and VLANs. Management and operation of failover and link aggregation is handled once in the networking fabric, in contrast to a general-purpose operating system like Windows Server which requires that you configure every server.

Microsoft Hyper-V

With Windows Server 2012, Microsoft introduced the ability to team anywhere from one to 32 NICs. (A team created on a single NIC may be used to separate traffic using VLANs, but it does not provide any failure protection).

Prior to Windows Server 2012, Microsoft always relied on the NIC vendor to provide any teaming software. Microsoft's capability is implemented to be hardware agnostic. Microsoft has built automation capabilities around it with PowerShell and SCVMM, making it a popular choice for many Microsoft shops. If NIC teaming is deployed, specify *Hyper-V switch port* as the traffic distribution method for Hyper-V 2012. Improvements were made in Hyper-V 2012 R2, changing the recommendation to use *Dynamic* as the traffic distribution method.

Before using Microsoft's NIC teaming feature, read Microsoft's documents explaining the proper use and configuration of teams:

- Hyper-V 2012: <http://www.microsoft.com/en-ca/download/details.aspx?id=30160>
- Hyper-V 2012 R2: <http://www.microsoft.com/en-us/download/details.aspx?id=40319>

Cisco UCS

Cisco has always provided teaming-like capabilities in their fabric—the data plane of the fabric interconnects is always active/active. A single NIC defined on a physical server can provide failover capability because if one fabric interconnect fails, the surviving fabric interconnect takes over operations (assuming that the NIC was defined with fabric failover). Cisco UCS can provide NIC aggregation capabilities, since both fabric interconnects in the data plane normally operate in active/active mode.

In addition, the Cisco Fabric Interconnects can make use of port channels to connect to upstream Cisco Nexus Switches possibly configured as virtual port channels (vPC). vPC is a virtualization technology that presents two Cisco Nexus Switches paired devices as a single Layer 2 logical node to access layer devices or endpoints. A vPC allows links that are physically connected to two different Cisco Nexus devices to appear as a single port channel. Some benefits of vPC include:

- Uses all available uplink bandwidth
- Provides fast convergence upon link or device failure
- Simplifies network design
- Highly resilient and robust Layer 2 network
- Enables seamless virtual machine mobility and highly available server clusters

However, if the server is configured with a single virtual interface card (VIC, also known by the industry name as a converged network adapter, or CNA) and that VIC fails, recovery can occur if the server is part of a Failover Cluster. In that case, the loss of the NIC would cause the node to lose communication with the rest of the cluster. This would cause the services/resources on the server with the failed VIC to be restarted on another node of the Failover Cluster.

If the length of the service interruption caused by a restart on another node is not acceptable, a second VIC can be placed into the each server and multiple NICs can be created and teamed. When creating the NICs to be teamed, define one NIC of the team to exist on the first VIC to use one fabric, and do not enable failover. Define the second NIC of the team on the second VIC and to use the other fabric without failover enabled.

NIC teaming is often implemented to aggregate lower-speed NICs in order to gain throughput. Since Cisco UCS defaults to 10 Gigabit Ethernet connections, aggregation is generally not required. If it is required, you can team the NICs or consider using SR-IOV (VM-FEX), covered later in this document. NIC teaming and SR-IOV are mutually exclusive solutions: SR-IOV does not support teaming.

Quality of Service

Recommendation: For simple quality of service (QoS), define quality of service in the fabric. For cloud and datacenter environments, deploy the Cisco Nexus 1000V Series Switch because it provides consistent operational processes across physical and virtual environments and across hypervisors. QoS is just one of many aspects of network management handled by the network administrators. When you keep the management in the network components, you help network administrators avoid having to access to Windows Server utilities to define QoS. Additionally, the resources required to implement QoS are then performed on the hardware designed to handle this, instead of through a general-purpose operating system like Windows Server.

Quality of service (QoS) for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for critical applications that run on the network. By using QoS mechanisms, network administrators can plan and use existing resources efficiently to ensure the required level of service without resorting to over-provisioning or reactive expansion of networks.

Microsoft Hyper-V

Windows Server 2012 QoS includes new bandwidth management features that enable cloud hosting providers and enterprises to provide services to deliver predictable network performance to virtual machines on a server running Hyper-V. QoS provides bandwidth management, classification and tagging, priority-based flow control, policy-based QoS for the physical network, and policy-based QoS for the virtual network. To define and manage QoS settings, Microsoft employs a combination of PowerShell scripting and Group Policy management.

Cisco UCS and Nexus Switches

Cisco UCS and Cisco Nexus switches come with pre-defined classes of service that can be quickly defined and implemented for QoS using the tools the network administrators are familiar with using. Cisco QoS provides all the capabilities of those provided by Windows Server 2012. Additionally, the Cisco Nexus 1000V Series Switch (covered later) is available to enhance policies available.



Management IP Addresses

Recommendation: Use two Cisco UCS Manager Management IP Server pools when using SMI-S for server management; default ext-mgmt for non-SMI-S managed servers and one (or more) pools for SMI-S managed servers.

By default, Cisco UCS Manager provides for specifying out-of-band management IP addresses for each server managed by Cisco UCS Manager. This connection is to each server's Cisco Integrated Management Controller (CIMC). The default pool of available IP addresses for this purpose is defined in the **LAN > Pools > root > IP Pools > IP Pools** ext-mgmt. As new servers are discovered within the environment, an out-of-band IP address is assigned to the service providing for management access to that server through Cisco UCS Manager's CIMC. Servers that are not managed by Microsoft's SCVMM should use this IP management pool in order to obtain an out-of-band management IP address.

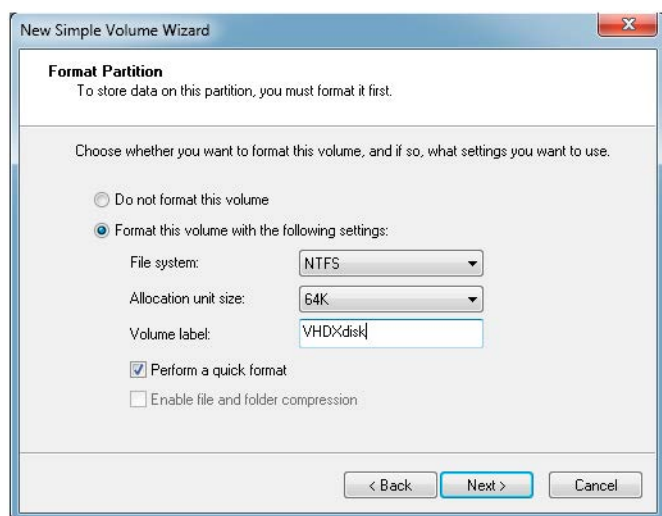
When servers are managed by Microsoft's SCVMM using the SMI-S protocol, it is recommended to provide a second management IP address pool. Servers managed by SMI-S should set their management IP address assigned from this pool and defined as an operational policy in a Service Profile.

Using a policy helps to ensure a single out-of-band management IP address stays with the Service Profile instead of with a specific physical server. For example, assume a server experiences a hardware failure that removes it from service. The Service Profile can be temporarily moved to another, similar server in order to continue operations. Any SMI-S operational policies or automation that was defined based on the IP address associated with the Service Profile would move to the new server. If the default ext-mgmt IP address were used, any SMI-S routine that was based on a management IP address would break because the Service Profile is now on a server with a different management IP address.

Host Storage

Hyper-V can use any Windows Server supported storage solution. This means that any form of Direct Attached Storage such as SAS and SATA, even USB drives, can be used to store VMs. Obviously, not all are good fits. Most Hyper-V installations are likely to be Failover Cluster installations, so shared storage becomes something that must be planned.

If the Hyper-V system is going to be a standalone, or nonclustered, system, your typical storage considerations must be evaluated. USB is not really appropriate for any production work. It works fine for transporting VMs from one machine to another if you don't want to use the bandwidth necessary to copy the machines across the network. But most of the time, you will probably be using some sort of RAID configuration on local storage. Storage used for virtual hard disks should be formatted with a 64K allocation unit size, as shown in the following screen capture:



In addition to direct attached storage (DAS), Microsoft supports FC, FCoE, iSCSI, Server Message Block (SMB) 3.0, and Microsoft Storage Spaces for storage of virtual hard disks. Once the storage is connected to the system, it is all treated the same, so you should format volumes with the 64K allocation unit size.

Prepare for Virtual Machine Storage Access

Recommendation: Fibre Channel over Ethernet (FCoE) access. The rate of technology improvement is occurring fastest in the networking space. 40 Gigabit Ethernet networks are available and 100 Gigabit Ethernet is not far behind. Building now on 10 Gigabit Ethernet networks lays the groundwork for the highest performing environments.

Over the years, Microsoft has increased the number of ways VMs can access storage. Initially, you could access a virtual hard drive that was a file on a drive owned by the Hyper-V host or an iSCSI connection to a LUN. Then Microsoft added the ability to access Fibre Channel storage with a virtual host bus adapter (HBA). And beginning in 2012, Microsoft added the ability to use SMB (network attached) storage. At the same time, the Hyper-V hosts can also use this same storage. Of course, for that to be possible, the Hyper-V host needs to be configured with the proper adapters to be available to the VMs.

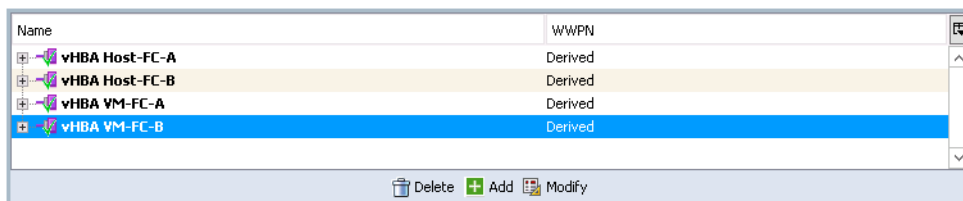
If you plan to provide direct access to storage from the VMs, it is recommended that you provide separate adapters for the virtual machines. Technically, Hyper-V allows for the storage adapters to be shared with the physical host, but it is recommended to keep them separate. For high availability, it is recommended that two adapters be defined, one on each fabric. The configuration will vary slightly depending on the type of storage access provided to the VMs.

The configuration can also vary depending on what the customer wants to do. It is possible for the host to access storage in one manner and the VMs to access it in another. For example, the host may be set up to access its storage for booting and storage over FCoE, but the customer may want to have the VMs access their storage via SMB. In that case, you would define FCoE for the host but not the VMs, and SMB for the VMs but not the host.

It is recommended that QoS be defined for accessing the storage, regardless of the protocol. Additionally, define the maximum transmission unit (MTU) for the QoS to be 9000 to enable jumbo frames on the network.

Fibre Channel over Ethernet and Fibre Channel

Fibre Channel access can be provided via traditional Fibre Channel (FC) components or via Fibre Channel over Ethernet (FCoE). The definition of virtual host bus adapter (vHBA) templates is accomplished in the same manner for FC as for FCoE. Generally, when providing FC access for VMs, the hosts will also be booting from FC. In this case, you should define two vHBA templates, one to run on each side of the fabric. It is recommended to define four vHBAs within your Service Profile template for the Hyper-V host, two for use by the host and two for use by the VMs. Two vHBAs would use the vHBA template for Fabric A and two vHBAs would use the vHBA template for Fabric B. Place the host's vHBAs before the VM vHBAs in the placement order, as follows:



Name	WWPN
vHBA Host-FC-A	Derived
vHBA Host-FC-B	Derived
vHBA VM-FC-A	Derived
vHBA VM-FC-B	Derived

Multipath software needs to be configured to ensure proper presentation of the storage via multiple paths to either the host or the VM. Microsoft's in-box MPIO works fine. Storage vendors have generally built their own multipath software on top of Microsoft's MPIO to enhance the management. Cisco makes no recommendation on which to use, Microsoft's or the storage vendor's; either works.

iSCSI

As with a FC SAN, it is recommended to have separate adapters for the host and VMs, two for the hosts and two for the VMs. In this case, the adapters are vNICs instead of vHBAs. But the same concepts exist. Define two vNIC templates: one for Fabric A and one for Fabric B. It is important that these templates be configured to run on only a single fabric, as MPIO will also be configured on this access.



Two separate network subnets should be created, one assigned to each Fabric. Two vNIC templates should be defined: one each for Fabric A and Fabric B. Do not enable failover on these NICs. The Service Profile template should then be defined with four iSCSI NICs. Two NICs would be on Fabric A and two NICs would be on Fabric B. As with FC/FCoE, MPIO must be configured for proper failover and aggregation.

Note: Do not build these NICs into teams. Microsoft does not support NIC teaming for iSCSI access.

SMB 3.0

SMB storage for Hyper-V virtual hard drive storage is new with Windows Server 2012. The SMB protocol has been around for a very long time, but its use as a storage protocol for clustering and applications that typically use block storage is new. As a result, you may not find many installations that are using this, but its ease of configuration may win over converts. With FC/FCoE and iSCSI, you need to configure access with World Wide Names (WWNs) or iSCSI Qualified Names (IQNs). Once access is configured, MPIO has to be configured. SMB simply needs access to a network share and it will automatically use up to 32 NICs to access the storage, automatically adapting to the addition or loss of NICs assigned to it.

From a configuration standpoint, it is recommended to configure this as you would configure iSCSI: a network subnet for each fabric and a vNIC template for each fabric, defined without failover. The Service Profile template using SMB storage for host and VMs should have four vNICs, with two to be used for the host and two to be used for the VMs. But, as noted above, there is no need to define MPIO. SMB automatically takes care of that.

SMB 3.0 is available as an option on the major storage vendors' SANs. It can also be provided by Microsoft Windows Server 2012, working as a file server using whatever storage it wants. This includes the Microsoft Storage Spaces. For SMB to be used for storage with Hyper-V, it must be Version 3.0 or later. Earlier versions, such as those provided by many NAS vendors and Samba implementations do not yet provide this latest version of SMB.

Host Clustering

Storage

Recommendation: Use FCoE for cluster storage. If customer SAN does not support FCoE connection, use FC. Once you've configured FCoE or FC to be available to the cluster, add the volumes to be used for storage of virtual hard drives to Cluster Shared Volumes.

Failover Clustering is a feature available on all server versions of Hyper-V. It is not available to Hyper-V when running the client version of Hyper-V. As Microsoft has added different storage protocol support for the host operating system, it has also added different protocol support for Failover Clustering. Currently, Microsoft supports FCoE, FC, iSCSI, SAS, and SMB 3.x as storage protocols for shared storage in a Failover Cluster. With its partnerships with EMC and NetApp, Cisco supports FCoE, FC, iSCSI, and SMB 3.x. Currently, Cisco does not have a solution for SAS shared storage.

NICs

The number of NICs and their assignment is quite dependent on the application being deployed. The number of NICs and their usage is not unique to Cisco UCS, but the flexibility provided by VICs allows for more control to be applied. Table 1 describes the types of NICs that are typically configured in a Hyper-V cluster.

Table 1. NICs Typically Configured in Hyper-V Clusters

Usage	Comments
Management	<ul style="list-style-type: none">• Only available to the Hyper-V parent partition• Used for cluster and host management• Ensure this network is configured to register its connection in DNS• Ensure this network is listed first in Network Binding Order• Allow cluster network communication on this network (one of at least two networks so configured)• Be sure Cisco UCS Manager failover is enabled

Usage	Comments
Live Migration	<ul style="list-style-type: none"> • Only available to the Hyper-V parent partition • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • Ensure that this NIC is selected in Live Migration Settings... • (Optional) Check <i>Allow cluster network communication on this network</i> • Define a QoS policy in Cisco UCS Manager and on Cisco Nexus • MTU = 9000 • If you plan to live migrate to/from nonclustered Hyper-V hosts, this network should be defined on the nonclustered hosts. • Be sure Cisco UCS Manager failover is enabled
Cluster Shared Volume	<ul style="list-style-type: none"> • Only available to the Hyper-V parent partition • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • Use PowerShell to set network prioritization to lowest priority (see next section) • Check <i>Allow cluster network communication on this network</i> (one of at least two networks so configured) • Ensure <i>Client for Microsoft Networks</i> and <i>File and Printer Sharing for Microsoft Networks</i> are enabled to support Server Message Block (SMB), a requirement for CSV • MTU = 9000 • Be sure Cisco UCS Manager failover is enabled
(Optional) iSCSI (2/4)	<ul style="list-style-type: none"> • Do not use NIC teaming • Place NICs on separate subnets • Cisco UCS Manager placed on separate fabrics – failover disabled • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • If using for host storage, uncheck <i>Allow cluster network communication on this network</i> • Disable <i>File Sharing and Printer Sharing</i> • Define a QoS policy in Cisco UCS Manager and on Cisco Nexus • MTU = 9000 • Configure for multipath I/O using either Microsoft's MPIO or the storage vendor's DSM • Use separate NICs for host and guest iSCSI access
(Optional) SMB 3.0 (2/4)	<ul style="list-style-type: none"> • Configuration of NICs (number, subnets, failover) is dependent on capabilities provided by SMB storage provider • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • If using for host storage, uncheck <i>Allow cluster network communication on this network</i> • Consider defining a QoS policy in Cisco UCS Manager and on Cisco Nexus • MTU = 9000 (dependent on SMB storage provider) • Do not configure MPIO; it uses its built-in multichannel I/O • Use separate NICs for host and guest access
VM NICs	<ul style="list-style-type: none"> • Similar configuration to above. Covered in a subsequent section on Virtual Machine Clustering.

Network Prioritization

By default, all internal Hyper-V Failover Cluster networks have a metric value. The starting value for this has varied with different releases of the operating system, but the hierarchy of the values has remained a constant.

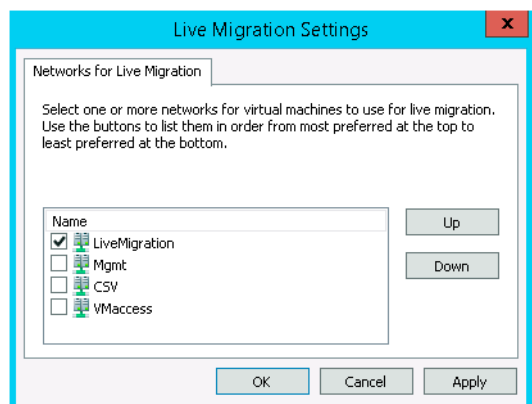


When you create Cluster Shared Volumes on a Failover Cluster, the clustering software selects a network that appears to be the best for CSV traffic and assigns the lowest metric value to that network. Often this is not the network that you may have configured for CSV communications. The second lowest network designates the network to be used for Live Migrations. To ensure the network you configured for CSV is used for CSV, use the following PowerShell commands. (CSVnetworkname is the name you gave to the CSV network).

```
Get-ClusterNetwork "CSVnetworkname" ).Metric = 900
Get-ClusterNetwork | FT -auto Name, Metric
```

The value of 900 is lower than any default metric assignment, ensuring that the CSV network will have the lowest value.

The Failover Cluster Manager console provides an easy way to designate your Live Migration network. Select Live Migration settings from the Actions pane of Failover Cluster Manager console. Select the network(s) you have configured for Live Migration usage. If you have multiple selections, you can use the Up and Down buttons to adjust the order to be used for Live Migration, as follows:



Host Memory

Host Memory Calculation

Recommendation: Ensure that you allow at least 2 to 8 GB for the host, depending on the number of VMs running. (This recommendation is not unique to Cisco UCS). The host needs memory to run the Hyper-V parent partition and the various management services. The actual amount varies, depending on the number of VMs running and management features you might run in the parent partition.

In order to understand the memory needs for the host, it is necessary to understand a bit about the architecture of Hyper-V. Hyper-V consists of two different components. The hypervisor, that portion that owns the physical hardware, is a very tightly written piece of code that takes less than 1 MB of memory. This is the piece that owns and schedules the CPUs and memory on the system. In addition to the hypervisor, Microsoft implements what they call a “parent partition.” This could be considered a special-function virtual machine. It is this parent partition that you log into when you log into any platform that has Hyper-V installed. This is the area that handles the interfaces between virtual NICs and physical NICs and virtual HBAs and physical HBAs. It is also the location in which the device drivers are installed, as well as other management capabilities like clustering, Hyper-V Manager console, and Windows Management Instrumentation (WMI), and so on.

Some people advise that the host needs only 512 MB or 1 GB of physical RAM for these management components and the rest can be used for VMs. Others advise that it is good practice to allow 2 GB of RAM for the host and the rest can be used by VMs. Neither approach tells the whole story.

Each VM also requires some workspace in the host's memory, above and beyond what the host is using. If you create a VM with 1 GB of memory or less, there will be about 40 MB of memory overhead allocated at the physical level in addition to the 1 GB (or less) allocated for the VM itself. As the amount of memory in the VM increases, it will require another 8 MB of memory overhead for each 1 GB increment of memory in the VM. The formula for calculating this is $32 \text{ MB} + ((\text{VM Memory} - 1) * 8 \text{ MB})$.

Table 2 shows some examples.

Table 2.

Amount of Memory in the VM	Amount of Memory Overhead in the Host
Up to 1 GB	32 MB
2 GB	40 MB
4 GB	56 MB
8 GB	88 MB

The observant reader will see a discrepancy between the entry for 1 GB and the amount stated above as required for the 1 GB VM. Above it says 40 MB is required for the 1 GB VM, whereas in the table it shows 32 MB. What's the difference? A management process also runs in the parent partition for each actively running VM. This management process requires about 7 to 8 MB of host memory. So, in addition to determining memory overhead, there is an additional 8 MB that needs to be added for each running VM.

This does not sound like a lot of memory, but if you are going to run 100 2-GB VMs, in addition to the 200 GB of memory required by the VMs themselves, you will need to add 4.8 GB for memory overhead and management processes. Additionally, this assumes that there are absolutely no other services running, such as clustering, file services, print services, remote desktop access for management, antimalware agents, central management agents, and so on. It is not recommended to install other roles such as file sharing to a Hyper-V server, but there are almost always going to be other management components.

Paging File

Recommendation: For Windows Server 2012 Hyper-V and Windows Server 2012 R2 Hyper-V, the page file of the management OS (commonly called the host OS) should be left at the default setting of System managed size. This is per the [Hyper-V product group](#). For Windows Server 2008 Hyper-V and Windows Server 2008 R2 Hyper-V, reduce the size of the paging file on the system drive. (This recommendation is not unique to Cisco UCS).

A paging file is an area on the hard disk that Windows uses as if it were RAM. This is used for two primary reasons. One of the more critical ones with today's systems is to capture a memory dump if the system crashes. The second one has become less critical with today's systems: if there is not enough physical RAM to hold all the currently running processes, portions of the processes memory is paged out to the paging file, to be retrieved when it is needed.

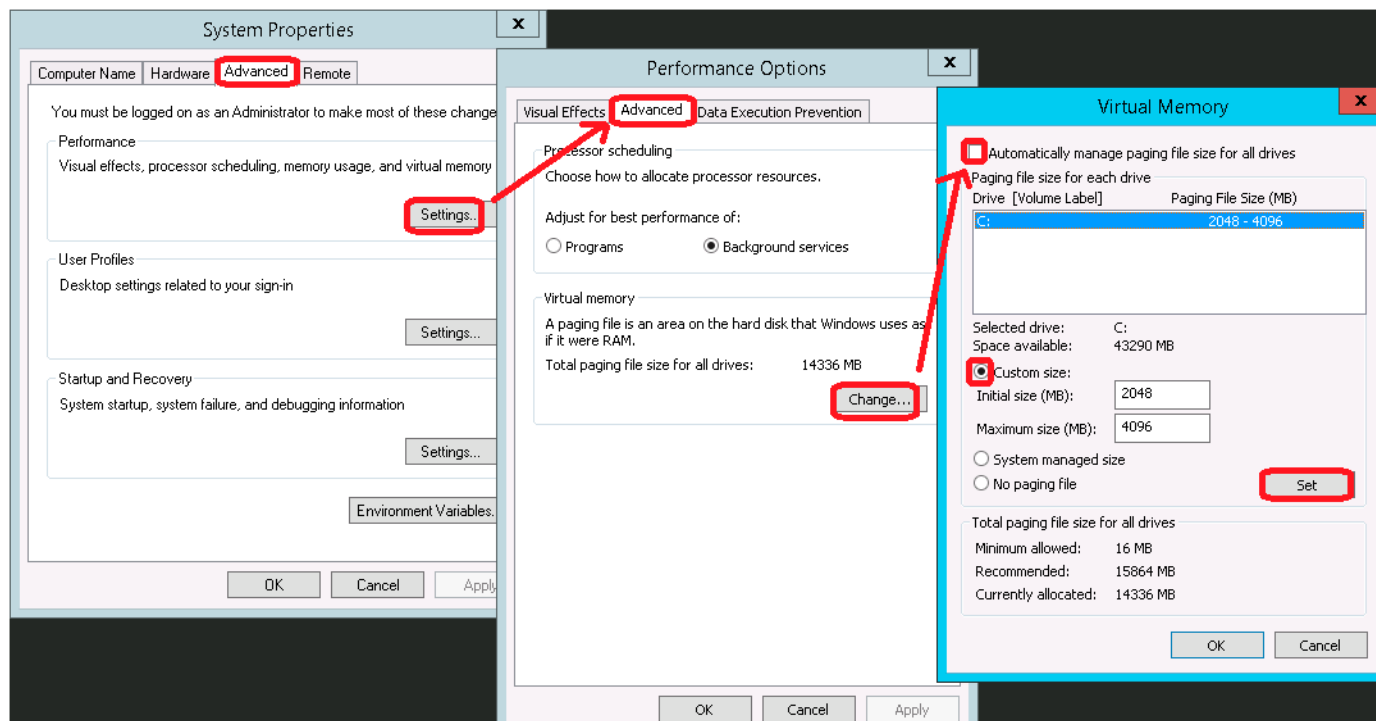
Hyper-V hosts tend to be configured with a significant amount of memory in order to host all the VMs. But Hyper-V itself does not require lots of RAM. And this is where the second usage of the page file becomes almost meaningless for today's servers. Yes, there are some services that still require the existence of a paging file, but if a Hyper-V host is configured properly, Hyper-V will most likely not need a paging file.

By default, Microsoft configures the paging file based on the physical RAM in the system. This tends to waste space on the system volume. It is recommended to minimize the size of the paging file on the system drive. Hyper-V can technically run in about 1 GB of memory. It uses additional memory as table space, which will grow with the number of virtual machines running. But rarely will it need more than 2 GB of RAM for its own use. The rest of the physical memory is assigned to the VMs. It is a recommended practice to keep some space for the paging file on the system drive in order to capture a partial memory dump should the system crash. A setting of 2 to 4 GB should be sufficient.

The following screen captures show the process. Looking at the last window, you can see that the system automatically assigned a paging file of 14 GB on a system with 96 GB of physical RAM. By changing the setting, you can save 10 to 12 GB of storage on your system drive, and more on systems with more memory.



In System Properties, select the *Advanced* tab. Click *Settings...* Again, select the *Advanced* tab and click *Change...* under Virtual Memory. Uncheck the *Automatically manage paging file size for all drives* and click on the *Custom size* radio button. Enter an Initial size of 2048. If you want to allow the paging file to grow beyond that size (probably not needed), enter a larger value in *Maximum size*. You must click the *Set* button to put the change into effect. This will require a reboot. The reboot will return the extra storage assigned to the paging file on the system drive for use for other purposes.



Virtual Machines

Virtual Machine Generation

Recommendation: Adopt the use of Generation 2 VMs cautiously and with appropriate knowledge of the current limits placed on them. Two areas are good candidates. Generation 2 VMs provides a higher level of VM security by implementing a secure boot process. Generation 2 VMs can boot from the standard network adapter instead of the legacy adapter, providing for a higher-speed Preboot Execution Environment (PXE) boot. If VMs are regularly created using PXE, the process will go faster and put less overhead on Hyper-V than the legacy adapter required hardware emulation.

In 2012 R2, Microsoft introduced a new generation of virtual machines. The generation determines the virtual hardware and functions presented to the virtual machine. Because different hardware is presented, there are limitations in the portability of the VM and the operating system environments that can run within the VM. Once a VM is created as one generation or the other, it cannot be converted to the other generation; it would need to be rebuilt to the other generation.

A Generation 1 VM provides virtual hardware to the VM that can be used in any version of Hyper-V. This is important if a Hyper-V earlier than 2012 R2 continues to host VMs.

A Generation 2 VM provides virtual hardware to the VM that can be used only on Hyper-V 2012 R2 and later. IDE disk and legacy network adapter support has been removed. Only Windows Server 2012, Windows Server 2012 R2, and the 64-bit versions of Windows 8 and 8.1 can be the guest operating systems running in a Generation 2 VM.

Virtual Machine Networking

NICS

Recommendation: Create one or more Hyper-V External virtual switches that are accessible only to the virtual machines (See Table 3). Note that NIC recommendations for VM Failover Clusters are covered in a later section.

Table 3.

Usage	Comments
VM Access (this could be a normal Hyper-V NIC or it could be a VEM if the Cisco Nexus 1000V is deployed)	<ul style="list-style-type: none">• Only available to the virtual machines• Ensure this network is configured to register its connection in DNS• Ensure this network is listed first in Network Binding Order• Cisco UCS Manager failover enabled

Cisco Nexus 1000V

Recommendation: Deploy the Cisco Nexus 1000V for any Hyper-V 2012 or later deployment that is managed by SCVMM. The Essential version of the Cisco Nexus 1000V Series Switch provides base network management capabilities such as VLANs, private virtual LANs (PVLANS), ACLs, QoS, Link Aggregation Control Protocol (LACP), and multicast. More sophisticated features like Dynamic Host Configuration Protocol (DHCP) snooping, IP source guard, dynamic ARP inspection, and Cisco VSG require the Advanced version of the Cisco Nexus 1000V

Prior to Hyper-V 2012, there were no options available to change Hyper-V's virtual switch or how it was managed. In essence, virtual switches had to be individually managed from each host. Hyper-V 2012 changes that with the implementation an extensible virtual switch, allowing third-party vendors to extend the capabilities of Hyper-V's basic virtual switch. The Cisco Nexus 1000V is such an extension, creating a distributed virtual switching platform with advanced networking features, integrated virtual services, and a consistent operational model across physical and virtual environments.

The Cisco Nexus 1000V Series Switch is implemented as a pair of virtual machines running the NX-OS operating system. Although the switch it can be deployed as a single VM, it is recommended to always deploy it as a highly available solution with two VMs. The VMs run the supervisory module (VSM), the management component that controls multiple Cisco Nexus Virtual Ethernet Modules. The Cisco Virtual Ethernet Module (VEM) is a software component deployed on each Hyper-V host as a forwarding extension to Hyper-V's extensible switch.

Management and control of the Cisco Nexus 1000V in a Hyper-V environment requires the use of SCVMM. All configuration policies defined on the VSM are automatically propagated to SCVMM to allow the SCVMM administrator to use these policies when creating virtual machines.

Cisco Nexus 1000V offers advanced networking features to Microsoft Hyper-V environments, including:

- **Advanced switching:** As private virtual LANs (PVLANS), quality of service (QoS), access control lists (ACLs), port security, and Cisco vPath.
- **Security:** Dynamic Host Configuration Protocol (DHCP) snooping, Dynamic Address Resolution Protocol (ARP) Inspection, and IP source guard to ensure that the Cisco Nexus 1000V Series is fully aware of all server virtualization events, such as live migration.
- Network services using Cisco vPath—for example, Cisco Virtual Security Gateway (VSG).
- **Provisioning:** Port profiles to address the dynamic nature of server virtualization, which enables you to define virtual machine network policies for different types or classes of virtual machines from the VSM and then apply the profiles to individual virtual machine virtual NICs (vNICs) through the Microsoft System Center Virtual Machine Manager for transparent provisioning of network resources. Port profiles are a scalable mechanism for configuring networks with large numbers of virtual machines.

- **Visibility:** Virtual machines participating in traffic monitoring activities, such as Cisco NetFlow or Cisco Encapsulated Remote Switched Port Analyzer (ERSPAN), which can continue these activities uninterrupted by live migration operations. With the ability to migrate network and security policies through live migration, regulatory compliance is much easier to enforce with the Cisco Nexus 1000V Series, because the security policy is defined in the same way as physical servers and constantly enforced by the switch.
- **Manageability:** Simple Network Management Protocol (SNMP), NetConf, syslog, and advanced troubleshooting command-line interface (CLI) features.

SR-IOV versus VM-FEX

Recommendation: Evaluate the specific customer situation. SR-IOV is a niche solution. Not all workloads benefit.

Single Root I/O Virtualization (SR-IOV) is an industry standard to assist when virtualizing demanding workloads that require high network performance. This is a technology that Microsoft fully supports in the Hyper-V environment, supporting even Live Migration of SR-IOV configured VMs between hosts with or without SR-IOV, but it is up to the hardware vendors to be able to provide SR-IOV on their adapters. Cisco Virtual Machine Fabric Extender (VM-FEX) is Cisco's implementation.

Performance is achieved by bypassing Hyper-V's extensible virtual switch. Microsoft has found lower CPU utilization for network functions, lower network latency, and higher network throughput by using SR-IOV within workloads that benefit from this technology.

Not all workloads benefit from this technology. For example, in general-purpose virtualized environments that share pools of compute, storage, and network resources running varying workloads, the ability to squeeze the absolute maximum performance out of each element does not have as high a priority as having flexibility in enforcing QoS policies. Microsoft QoS is enforced in their virtual switch, and Cisco enforces QoS in their Nexus switches. Use of Cisco VM-FEX and the Cisco Nexus 1000V are mutually exclusive. You cannot use the Cisco Nexus 1000V switch for traffic that is to be controlled by VM-FEX.

However, in situations that demand the most from a system, VM-FEX can have significant benefits. Using VM-FEX in a VDI environment can free up CPU resources that can be used to run more virtual desktop images, possibly allowing fewer servers to be deployed to host the needed number of VDI instances.

VM-FEX is not available for Hyper-V 2008 or 2008 R2. It can be configured manually on Hyper-V 2012. For use with 2012 R2, it requires SCVMM for management purposes.

Virtual Machine Storage

Recommendation: Assign VM storage to .vhdx virtual hard drives whenever possible. VHDX is compatible with Windows Server 2012, Windows 8, and later. If backward compatibility is needed, use .vhd virtual hard drives. Use of vHBA, iSCSI, or pass-through disks puts certain limits on the functions that can be performed on the VM, so their usage is discouraged. Use of VHDX or VHD for VMs helps to ensure maximum usefulness of the virtualization technology.

Hyper-V VMs have access to a variety of storage protocols. Up until Windows Server 2012 R2, it was a requirement for the boot volume to be connected to the VM via the IDE controller. This did not impact performance: it was just the way Hyper-V worked. With the release 2012 R2, Microsoft added the ability to boot from a SCSI controller. To boot from the iSCSI controller requires that the VM be created as a Generation 2 VM. Prior to Hyper-V 2012 R2, no version of Hyper-V can support a Generation 2 VM.

For data storage, IDE and SCSI access to virtual hard drives have always been supported. For these VM data drives, the virtual hard disk could reside on any disk storage device recognized by the Hyper-V parent partition, which was basically any storage available to the Windows Server environment. 2008 and 2008 R2 could also use iSCSI to present iSCSI LUNs to VMs. Windows Server 2012 introduced support for virtual FC adapters to present FC and FCoE LUNs to VMs running Windows Server 2008 and later. Microsoft introduced SMB 3.0 as a highly optimized network file storage solution for access by VMs running Windows Server 2012 and later and Windows 8 and later.

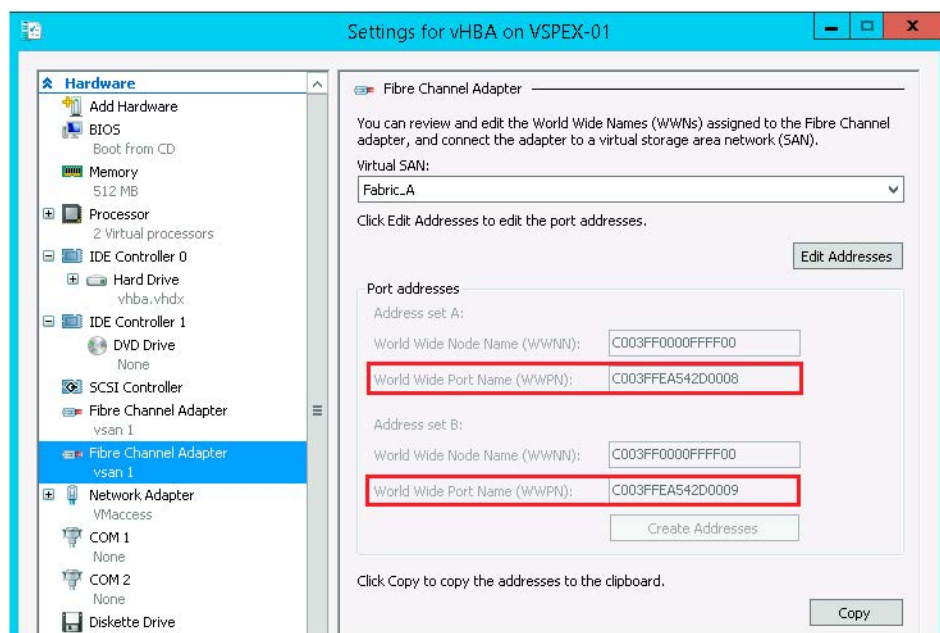
Hyper-V has always supported a device that is connected as a direct-attached physical disk on the Hyper-V host, also known as a pass-through, or raw, disk. This configuration would assign a VM exclusive ownership and use of a device connected to the parent partition. This configuration is never recommended—it restricts a VM to the particular server on which the physical device is attached, removing many of the benefits of a VM because it cannot be moved to any other host. In the earliest version of Hyper-V, there was a performance benefit to connecting a VM directly to a physical disk. Through improvements in subsequent releases of Hyper-V, this is no longer the case.

Virtual Fibre Channel

Recommendation: Use .vhdx files whenever possible. But if you use virtual FC, configure vHBAs as you do physical HBAs, with redundancy and multipathing software for high availability and link aggregation. Configure two vHBAs on the Hyper-V host for exclusive use by the VMs, one for each side of the fabric.

Microsoft introduced the ability to use virtual Fibre Channel adapters within VMs in 2012. This allows VMs direct access to FC storage LUNs for data volumes or for shared storage in a clustered environment. vHBAs can be created for use on either FC or FCoE. The ability to create vHBAs is dependent on the HBA driver supporting N Port ID Virtualization (NPIV). Cisco introduced support of NPIV on VICs with Cisco UCS Version 2.1.2.

Creation and use of vHBAs on FCoE for VMs is slightly different than for physical machines. Two NPIV WWPNs, or Hyper-V address sets, will be associated with each virtual FC adapter.



Both WWPNs must be zoned, masked, and/or registered to the appropriate storage in order for live migration to work for that virtual machine. When powered on, only one of the two WWPN addresses will be used by the guest at a specified time. If you request a live migration, Hyper-V will use the inactive address to log into the storage array and ensure connectivity before the live migration continues. After the live migration, the previously active WWPN will become inactive. It is important to validate connectivity and live migration functionality prior to putting a virtual machine into production using the virtual Fibre Channel feature.

Virtual Machine Clustering

Shared Storage

Recommendation: Used shared .vhdx file for the shared storage requirement in VM Failover Clustering. Shared .vhdx files must reside on a Cluster Shared Volume.

VM Failover Clusters can use the access methods shown in Table 4 to access shared storage: FCoE, FC, iSCSI, SMB 3.0, or a .vhdx file from a Cluster Shared Volume. By far the easiest to configure is the shared .vhdx file. Configuration of a shared .vhdx file is simply a checkbox on each VM that is part of the Failover Cluster. The next simplest configuration is SMB 3.0, because it requires just configuring a pair of NICs. FCoE and FC require configuring eight paths to shared storage, plus configuring MPIO. iSCSI is similar to SMB, but its configuration is more complex, and it also requires MPIO configuration.

Table 4.

Usage	Comments
VM access (this could be a normal Hyper-V NIC or it could be a VEM if the Cisco Nexus 1000V is deployed)	<ul style="list-style-type: none"> • Only available to the virtual machines • Uncheck <i>Register this connection's addresses in DNS</i> • Ensure this network is listed first in Network Binding Order • Check <i>Allow client connect through this network</i> • Check <i>Allow cluster network communication on this network</i> (one of at least two networks so configured) • Be sure Cisco UCS Manager failover is enabled
Cluster Communications	<ul style="list-style-type: none"> • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • Check <i>Allow cluster network communication on this network</i> (one of at least two networks so configured) • Be sure Cisco UCS Manager failover is enabled
(Optional) iSCSI (2)	<ul style="list-style-type: none"> • Do not use NIC teaming • Place NICs on separate subnets • Be sure that Cisco UCS Manager is placed on separate fabrics, with failover disabled • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • Uncheck <i>Allow cluster network communication on this network</i> • Disable File and Printer Sharing • Define a QoS policy in Cisco UCS Manager and on the Cisco Nexus 1000V • MTU = 9000 • Configure for multipath I/O using either Microsoft's MPIO or the storage vendor's DSM
(Optional) SMB 3.0 (2)	<ul style="list-style-type: none"> • Configuration of NICs (number, subnets, failover) is dependent on capabilities provided by SMB storage provider • Uncheck <i>Register this connection's addresses in DNS</i> • Uncheck <i>Allow client connect through this network</i> • Uncheck <i>Allow cluster network communication on this network</i> • Consider defining a QoS policy in Cisco UCS Manager and on the Cisco Nexus 1000V • MTU = 9000 (dependent on SMB storage provider) • Do not configure MPIO— it uses its built-in multichannel I/O

Summary

Hyper-V 2012 provided a tremendous step forward in server virtualization—and as a cloud platform—with enhancements in the Windows Server operating system and many enhancements in the Hyper-V role. The Cisco UCS platform combined with the Cisco Nexus 1000V Series Switch and Cisco VM-FEX technology provides the most complete platform for Microsoft Hyper-V, enabling organizations to not only take full advantage of the Hyper-V capabilities, but also to extend them through Cisco innovations. When PowerShell and System Center integration is used with Cisco UCS Manager, organizations gain a single, unified way to manage all aspects of the Microsoft and Cisco solution. It becomes very clear that together Microsoft Hyper-V and Cisco UCS provide the best and most complete solution in the market, enabling organizations to embrace virtualization and the private cloud.

Efficient management of the computing infrastructure necessary for a private cloud is very important for IT organizations. Cisco and Microsoft work together to provide an integrated management experience for workloads and scenarios, both physical and virtual. To that end, Cisco has developed comprehensive infrastructure management tools for your Cisco UCS data center—tools that work with and extend the ones you may already use to monitor, provision, configure, and orchestrate your Microsoft server and application software. To achieve this unified operations approach, Cisco has:

- Developed the Cisco UCS Management Pack for Microsoft System Center Operations Manager to help IT staff monitor health status for 1-to-N Cisco UCS platforms
- Developed the Cisco UCS Integration Pack for Microsoft System Center Orchestrator to help administrators automate and standardize Cisco UCS management
- Developed Cisco UCS integration for Microsoft System Center Virtual Machine Manager, comprising three components:
 - A UI extension that enables access to Cisco UCS within SCVMM
 - Management of Cisco Nexus 1000V Series Switches
 - Management of Cisco VM-FEX technology (coming soon)
- Integrated Cisco UCS with Microsoft Windows PowerShell through Cisco UCS PowerTool to make it easier for administrators to manage infrastructure alongside operating systems and applications on Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack-Mount Servers, and standalone Cisco UCS C-Series servers by using familiar PowerShell cmdlets

For more information on Microsoft integration, visit: www.cisco.com/go/microsoft

For more information Cisco Nexus 1000V Series, visit: www.cisco.com/go/nexus1000v

For more information Cisco VM-FEX technology, visit: www.cisco.com/go/vmfex

For organizations interested in implementing a fully integrated and proven architecture that uses the Cisco UCS and Microsoft Windows Server capabilities, Cisco partners with both NetApp and EMC to offer complete converged infrastructure solutions in the form of the FlexPod and VSPEX solutions, respectively. Having multiple storage options allows partners and customers to have the most choices when looking for integrated Microsoft Private Cloud solutions. FlexPod with Microsoft Private Cloud is an integrated NetApp and Cisco reference implementation of Microsoft Private Cloud Fast Track reference architecture, which delivers IT as a service (ITaaS) through a cost-effective, flexible, highly manageable, and automated infrastructure that will grow with your business. It is a combined infrastructure stack of NetApp, Microsoft, and Cisco technologies that is available through channel partners and is validated by Microsoft through the Microsoft Private Cloud Fast Track program.

FlexPod with Microsoft Private Cloud has the flexibility to be sized and optimized to accommodate many different use cases, including application workloads such as Microsoft SQL Server, Exchange Server, SharePoint Server, and others. FlexPod with Microsoft Private Cloud enables simplified management with repeatable deployments.

For more information, visit: www.cisco.com/go/flexpod

VSPEX Private Cloud for Microsoft Windows Server with Hyper-V is a complete virtualization solution, proven by Cisco and EMC to run Microsoft applications and delivered to you by your trusted partner. Designed for flexibility and validated to help facilitate interoperability and fast deployment, VSPEX works with your Microsoft-based environment while removing the complexity and risk that typically accompanies the design, integration, and deployment of a best-in-class solution. As with the NetApp solution, VSPEX has the flexibility to be sized and optimized to accommodate many different use cases.

For more information, visit: www.cisco.com/go/vspex

About the Author

Tim Cerling, Technical Marketing Engineer, Cisco

Tim Cerling focuses on delivering customer-driven solutions on Microsoft Hyper-V and System Center products within the Cisco Data Center Group. Tim has been in the IT business since 1979. He started working with Windows NT 3.5 on the DEC Alpha product line during his 19-year tenure with DEC, and he has continued working with Windows Server technologies since then with Compaq, Microsoft, and now Cisco. During his 12 years as a Windows Server specialist at Microsoft, he co-authored a book on Microsoft virtualization technologies, *Mastering Microsoft Virtualization*. Tim holds a BA in Computer Science from the University of Iowa.

Acknowledgements

For their support and contribution to the input and review of this Cisco white paper, we would like to thank:

Frank Cicalese: Technical Solutions Architect, Cisco

Jeff Silberman: Technical Marketing Engineer, Cisco

Don Stanwyck: Product Manager, Microsoft

Manjnath Ajjampur: Principal Datacenter Technologist, Microsoft

Roger Osborne: Premier Field Engineer, Microsoft



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

C11-730725-00 01/14