

Building or Extending a Hyper-V R2 Cluster

This document provides exact steps to establish a functioning node in a new or previously existing cluster of servers running Microsoft's Hyper-V R2 SP1. The steps are built around the specific hardware used in the construction of this paper, but should be generic enough to facilitate similar implementation on any other hardware capable of running the Hyper-V R2 SP1 product.

Product and company names contained in this document may be trademarks of their respective owners. The author of this document has no connection with any of the mentioned companies beyond the producer/consumer relationship.

This information is provided in the hopes that it will be useful and every effort has been made to ensure its accuracy. No warranty or guarantee exists that it is completely error-free or that it is suitable to any particular scenario or environment. Use at your own risk.



This work is licensed under a [Creative Commons Attribution-ShareAlike 3.0 United States License](http://creativecommons.org/licenses/by-sa/3.0/).

By Eric Siron

Original document: 12/15/2010

Latest revision: 7/14/2011

Based on a work at <http://etechgoodness.wordpress.com/>

Table of Contents

Introduction	3
Document Conventions	3
Windows Command Shell and Tools Notes	3
Reference Hardware and Software.....	4
Planning and Considerations	5
Memory.....	5
Physical Disks	5
Virtual Disks	6
Snapshots.....	7
Networking.....	7
Windows Server Operating System Virtualization Licensing Rights	8
Host Operating System Choices.....	9
Domain Controllers	9
Cluster Shared Volumes.....	10
Host Configuration.....	11
Pre-installation Hardware Configuration.....	11
Operating System Installation and Initial Configuration.....	11
Network Design.....	12
Initial Network Configuration	12
Enable Remote Access and Install Prerequisite Components	15
Advanced Network Configuration	16
Installing Dell OpenManage.....	17
Install Core Configurator and Finalize Network and Host Configuration	18
Equallogic Host Integration Toolkit Installation.....	19
Storage Setup for Hyper-V and Cluster Shared Volumes.....	20
Setup Hyper-V and Failover Clustering	22
Setup Hyper-V Virtual Networking	22
Creating or Joining a Cluster	23
Establishing and Using Cluster Shared Volumes.....	25
Final Cluster Validation	25

Networking Troubleshooting	26
General Problems Related to Virtual Switches	26
Recovering from Blue Screen Loop after Enabling Intel VMQ.....	27
System Configuration Manager Virtual Machine Manager (SCVMM).....	28
Installing the SCVMM Server	28
Installing an SCVMM Administration Station and Configuring the SCVMM Server	29
Wrapping Up	31
Override LiveMigration’s Automatic Network Selection	31
Special Settings for Virtualized Domain Controllers	31
Using Server Core Instead of Hyper-V	32
Windows License Activation for Server Core.....	32
Installation Worksheets and Checklists	1
Pre-Installation Planning.....	1
Pre-Installation Hardware Configuration and Initial Operating System Configuration	1
Network Configuration	2
Operating System Configuration.....	3
Software Installation.....	3
Hyper-V and Failover Clustering	3

Introduction

Document Conventions

Lines to be typed at a command prompt or console window will be presented as light green mono-spaced text in a dark box. All text will be upper-case except in instances in which case-sensitivity is necessary. Lines start with a greater-than symbol. Example:

```
> NETSH INTERFACE IP SHOW INTERFACES  
> PING 192.168.25.1
```

Variable fields that will require you to provide information will display the type of information needed in italics surrounded by angle brackets. Example:

```
> PING <<primary domain controller>>
```

Important notes will be presented in red text in a dashed box. Example:

ENSURE ALL EQUIPMENT IS CONNECTED TO A UPS.

Industry-defined “Best Practices” are followed whenever practical for a small or medium-sized business that is not constrained by heightened security needs, such as those as might be found at a financial or medical institution. Security is not ignored in this document, but neither is it allowed to obscure the underlying procedure being explained. If your organization requires these higher levels of security, have your security expert review the settings in this document, especially those related to the network and SAN configurations.

Hyper-V R2 SP1 was installed as the hypervisor for this document, but the instructions will indicate any differences that would be needed if you choose to use a 2008 R2 SP1 Server Core install. These directions could be used for a Full installation of Server 2008 R2 SP1, but no GUI directions are included.

Windows Command Shell and Tools Notes

- Using the TAB key at a command-prompt will invoke an auto-completion tool in some contexts, most notably file/folder names and PowerShell commands.
- The command shell is almost exclusively case-insensitive, but not all commands are. Notably, feature names passed to the DISM.EXE tool are case-sensitive.
- When working in a GUI environment, it is common to type “EXIT” to leave a command prompt. Within a Hyper-V or Server Core shell, this closes the only available command window. Press CTRL-ALT-DEL (CTRL-ALT-END if accessing in an RDP session) and press “Start Task Manager”. Go to File and select New Task (Run...). Enter “CMD” in the box and press OK. Instead of typing “EXIT” at the prompt, type “LOGOFF”.
- NETSH allows most of its commands and arguments to be abbreviated. For example, “IPV4” can usually be typed as just “IP”, “INTERFACE” as “INT”, etc. Care was taken to use minimal abbreviation in this document, but some shortened commands may still be present. Use abbreviation at your own risk.

Reference Hardware and Software

The following hardware and software was used in the creation of this document. If the hardware you use will be different, you will need to consult your hardware manufacturer's instructions in order to duplicate any hardware-specific settings presented in this work.

- > Hardware
 - Host: Dell PowerEdge R715, BIOS 1.4.1
 - Dual 12-core AMD 6176 SE
 - 4 onboard Broadcom 5709 NICs, driver 16.2.1, firmware 6.2.14
 - 4 PCI-E Intel Gigabit Server NICs, driver 12.5.5
 - Two 146GB 15k internal SCSI drives, mirrored in hardware, firmware D809
 - H700 RAID controller, firmware 12.10.1-001, driver 4.31.1.64
 - Enterprise iDRAC6, BIOS 1.70
 - 128GB RAM
 - Dell/Equallogic PS4000X, firmware 5.0.2, Windows HIT Kit 3.5.1
 - 2x Cisco SLM 2024, firmware 2.0.0.8
- > Software
 - Microsoft Hyper-V R2 with SP1
 - Microsoft System Center Virtual Machine Manager 2008 R2 SP1

Planning and Considerations

Memory

With the release of Service Pack 1 for Hyper-V R2, Microsoft introduced “Dynamic Memory”. This allows the hypervisor to adjust the amount of RAM available to individual VMs based on minimums and maximums that you specify. This technology can be used to assign more total RAM to VMs than is physically present in the host computer. In case of contention, the system can prioritize automatically or follow weighting parameters that you set.

While Dynamic Memory can increase VM density on hosts, it is not a technology that should be used lightly. There is absolutely no software technology that can make up for insufficient quantities of physical RAM. RAM starvation of a VM will crush its performance. Instead of over-relying on Dynamic Memory, purchase sufficient amounts of RAM to run all the VMs you anticipate using. If you will be converting existing physical servers to virtual machines, it’s likely that they were purchased with extra RAM “just in case”. That approach is unnecessary with VMs. Use Performance Monitor to track their memory usage throughout a representative time frame and determine how much RAM they actually need. In the virtualization conversion, reduce the RAM allocated to the VM accordingly.

Physical Disks

In order to use LiveMigration (the ability to move a VM from one Hyper-V to another without any down time for the VM), the VMs must reside on a physical location that is simultaneously available to all hosts, such as a SAN. If you will have a single host, then neither Live- nor QuickMigration will be an option, so you can just use internal storage.

The basic method of storage for the Migration paths is to establish an NTFS LUN on a SAN, then create a virtual machine’s VHD on that LUN. Only one host at a time owns that LUN, so when the VM goes through Quick- or LiveMigration, ownership of the LUN is transferred from the source host to the destination host. This will result in a momentary service interruption, but clients should not notice and no processes should be stopped. In this configuration, only one VM can exist per LUN; this is because of the singular ownership nature of NTFS volumes.

To solve the interruption and one-VM-per-LUN problems, Microsoft has introduced Cluster Shared Volumes. All hosts in a cluster will be able to access a CSV simultaneously, so there is no delay to LiveMigrate a VM and as many VMs can be placed on the LUN as it has capacity for. Check that your SAN supports CSV volumes; if it allows multiple connections to one LUN from different hosts, then it most likely does support CSV.

In addition to the .VHD files that contain what the virtual machines consider to be their hard drives, a .BIN and a .VSV file will be created for each guest. The .BIN file will store the RAM state of a saved machine, but it also always exists while the machine is powered on. The size of this file will match the amount of RAM dedicated to the guest. If the guest is using Dynamic Memory, the file size will match its current allocated size. The .BIN file is deleted if the machine is powered off. The .VSV file contains the active configuration for a guest in a saved state and will average around 21 megabytes in size. Like the

.BIN file, it is also pre-allocated so that it exists any time the guest is powered on. Neither of these is actively written to while the machine is in a normal powered on state, but they must exist the entire time it is on. This is to ensure that a Save State operation occurs as quickly as possible. There is currently no way to bypass this functionality, so all guests will always have these files, even if the guest is configured to always shut down or shut off. For a typical deployment of virtualized servers, this probably doesn't matter. If you will have a very large number of guests, more common in the case of a VDI deployment, you must ensure you have sufficient drive space to hold these additional files.

Virtual Disks

The Microsoft virtual drivers do not allow booting a VM to a virtual SCSI volume, because they require access to the VMBus driver which is not available until the OS is loaded. As a result, you'll need to create all boot volumes on virtual IDE drives. You'll also need to use IDE for any guest systems in which you cannot install Integration Services. Since these are virtual drives, their performance is going to be mostly the same as a virtual SCSI drive. The only realistic limitations of virtual IDE drives are that the guest VM can have no more than four IDE devices at a time, and, just like it would with a physical IDE controller, will only talk to one IDE device per controller at any one time. Virtual SCSI drives can be added in much greater number and can be hot-added/removed. The capacity limit for both types is 2 terabytes. After a disk is created, it can be moved from the IDE to the SCSI chain and vice versa.

When creating a virtual disk file (VHD), you can choose between pass-through, dynamic, or fixed. A pass-through disk directly connects the VM to a drive or LUN attached directly to the host the VM is running on. A pass-through disk is used when performance is critical, because it allows the VM to directly handle I/O without involving a virtualization layer. Because the VM is handling I/O by direct communication with a single host computer, LiveMigration becomes problematic. When a LiveMigration event occurs, pass-through disks must be dismounted from the source host while the VM moves to the other host, which must then be mounted by the destination host. This will definitely cause a pause in operations and will probably disconnect or timeout remote activities.

When performance is critical, it is recommended to use a fixed disk instead of a pass-through disk. Microsoft claims that a Hyper-V fixed disk will perform just as well as a pass-through disk, and it has none of the drawbacks of pass-through. Fixed disks are recommended for high-use SQL database servers and Exchange servers. In all other cases, use dynamic disks. When creating a fixed disk, all of the requested space is allocated on the target LUN immediately. Dynamic disks allocate new space as needed and retain it until compacted. The overhead involved in expanding the dynamic disk and in tracking its variable state have a minimal but measurable performance hit. Dynamic disks can be converted to fixed disks and vice-versa. Either type can be expanded up to the maximum of 2 TB. It is possible to reduce the capacity of a VHD, but it requires a third-party tool such as "VHD Resizer".

There is a fourth VHD type call a differencing disk. It is only used with snapshots (checkpoints) and will be discussed further in the next section.

Snapshots

Snapshots utilize a fourth type of VHD file called the differencing disk. When a snapshot is created (called a checkpoint in System Center Virtual Machine Manager), Hyper-V creates a copy of the VM's configuration file and stops writing to its VHD(s). It creates an AVHD file for each VHD in the VM and writes changes there. The AVHD files are the differencing disks. When a file is read in the VM, the host examines the differencing disk to see if it needs to read from it or the base VHD. It is possible for multiple snapshots to be taken and for entire trees of snapshots to exist. They are useful when making changes to the VM, such as patches or even OS upgrades, but they are not useful for backup purposes. Differencing disks cannot be separated from their base VHDs and there is no way to selectively restore from a snapshot. Most importantly, active snapshots can rapidly consume large quantities of disk space. Especially for this reason, it is advisable to be judicious in snapshot usage. During the planning phase, assign sufficient disk space to handle the expected VM load as well as snapshots. Hyper-V Manager allows you to designate a special location for the snapshots of an individual VM, but not globally for a host or VM group.

Networking

Networking is probably the most complicated aspect of a Hyper-V deployment and should receive the most time and attention in the planning phases.

For optimal performance in a Hyper-V cluster, at least 4 NICs are required. The first will be the management connection for the host, the second will handle VM traffic, the third will be used for LiveMigration, and the fourth will carry cluster and CSV data. Microsoft supports using the VM NIC for management functions, but with some NICs, this can cause problems that cause your host to become unreachable and in need of a fairly involved repair at a late phase of configuration. For that reason and to avoid a performance hit, use a dedicated NIC for host management. For the best network configuration, add another adapter to split cluster from CSV traffic. Both lines will be mostly dormant in most cases, but traffic segregation can ensure they will never bottleneck each other. If you expect to have a moderate or high number of VMs, then adding NICs for VM traffic will be helpful. If you will be connecting to an iSCSI SAN, at least one NIC will be required for that as well.

Microsoft's Failover Cluster technology requires that no NIC that will participate in any type of cluster communication share a subnet with any other NIC on the same host. You will be required to use one subnet each for management, cluster communications, CSV traffic, and LiveMigration. Although it won't be part of cluster communications, you should also have iSCSI traffic in its own network. If you follow recommendations and do not allow the NICs for your management connections to be used for anything else, then they won't even have IP addresses of their own. If you do allow them to participate, they must exist in a separate subnet as well.

It is not recommended that you use teaming for anything. There are no RFCs that govern that technology, so each manufacturer is free to implement it in any way that they wish. There are seemingly unending quantities of reports of failures that appear to be related to teaming NICs in a Hyper-V environment. Use MPIO for your iSCSI connections and establish one virtual switch per NIC to be used in

VM traffic. None of the other functions will generate enough traffic on a regular basis to gain enough benefits from teaming to offset the risks.

The NIC(s) used for VM traffic will be converted into a virtual switch by Hyper-V. This is a well-featured layer-2 switch and fully supports VLAN trunking. When a VM is attached to a network created on it, a virtual switchport is created for it. That switchport can be assigned a VLAN, and the VM will then be a member of that VLAN. Multiple virtual NICs can be assigned to a single VM, which will allow it to participate in multiple VLANs simultaneously. If you have multiple NICs for virtual switches, and therefore multiple virtual networks, a single VM can have NICs connected to different virtual switches. The virtual NICs cannot be teamed.

Hyper-V does support advanced network features such as VMQ and VM Chimneying, but it is dubious just how much of a performance boost they actually add and are often problematic. Even though a NIC may support one or more of these technologies, it may be impossible to enable it in the host OS. Enabling the VMQ on an Intel card will cause the host to halt with a blue screen IRQL_NOT_LESS_OR_EQUAL error, which requires an involved clean-up process.

Windows Server Operating System Virtualization Licensing Rights

Each edition of Windows Server grants a specific number of licenses to be used in virtualization in addition to its normal physical unit license. Standard Edition allows you to install Server 2008 R2 on one physical host and install a single copy of Server 2008 R2 in a VM on that host. Enterprise Edition allows you to install Server 2008 R2 on one physical host and in up to four VMs on that host. Datacenter Edition allows you to install Server 2008 R2 on a per-processor basis. You can run an unlimited number of virtualized installations of Server on a licensed processor. Each edition allows you to run the same or a lower edition in the VM than on the host (i.e. a Standard Edition VM on an Enterprise Edition host). They also allow you to run previous versions of Server in your VMs in place of 2008 R2 (some restrictions on downgrade may apply if you use an OEM edition of Windows Server; check with your vendor). The virtual licenses are granted regardless of the hypervisor that you select. If you purchase a copy of Windows Server 2008 R2 Enterprise Edition and designate it for a particular host and use Hyper-V R2 or VMWare or Xen as the hypervisor, you still have rights to install four copies of Windows Server in VMs on that host. The physical license for your Enterprise Edition is still bound to that server; you cannot assign it on another physical host, even if you won't be using your virtualization rights there.

Technologies such as LiveMigration and vMotion can unintentionally bring you out of compliance when using virtualization rights. For instance, imagine you assign two Enterprise Edition licenses to two physical hosts in a cluster and use their virtualization rights to install four guest VMs on each machine. For maintenance purposes, you transition some of the VMs from one host to the other. For the duration of time that host is running more than four VMs, you are out of license compliance. However, in the event of a catastrophic failure of one host that would cause an automatic LiveMigration or vMotion to the other, the host license of the physical failed machine could be assigned to the surviving host, provided that it's not an OEM license. Microsoft's Failover Clustering tool and System Center Virtual Machine Manager can help you assign guests to hosts to make them somewhat "sticky" and design automated failover/failback triggers to help keep you in compliance. You can also, of course, license

individual VMs as though they were physical installations, and they can then run on any one host at any time and remain in compliance.

Host Operating System Choices

You can choose to install Hyper-V natively or select any of the Windows Server 2008 R2 editions except Web Edition. This guide is geared toward the native Hyper-V installation, but a Server Core installation would be nearly identical. The positives for installing Hyper-V natively are: it needs no license, it is the slimmest possible Hyper-V installation, and as long as Microsoft chooses to offer Hyper-V at no cost, you can continue upgrading it. Since Hyper-V is actually a version of Server Core with many roles and features removed, you can install applications and drivers on it that would install on Server Core.

If you choose to use a fully licensed edition of Windows Server, keep in mind that Standard Edition does not have access to the Failover Cluster feature, so it will not have access to LiveMigration or any other high availability features. A full installation of Windows, as opposed to a Core installation, is easier to work with and is not much “heavier” in terms of resource usage than a Core install, but it does have a greater attack surface and will need more patches. Any usage of any level of Windows Server will provide a greater attack surface and higher resource utilization than installing Hyper-V natively. Server Core does not allow the usage of RemoteFX, but native Hyper-V and a full Server installation do.

Domain Controllers

A common question is whether or not a domain controller should be virtualized. Many factors prevent the existence of a single correct answer. If your Hyper-V deployment is clustered, it is generally a good idea to have at least one virtualized DC. The primary danger is the Save State process. If a domain controller ever has its state saved, you must discard the saved state before turning it back on. If this is not done (and an automated process could take the decision away from you), the domain controller may stop participating in replication, and even in the best case, that will cause you a lot of headaches. If you do build a domain controller VM, ensure you set its Automatic Stop Action to “Shut Down”

A second concern is that if all your domain controllers are virtualized and you ever completely shut down your Hyper-V environment, your Hyper-V server(s) will have to be brought online without benefit of a domain controller. In general, this should not be an issue because of cached credentials and because most vital services on a Hyper-V host will run in the Local System context. A related tertiary concern is that if your Hyper-V cluster is crashed or unreachable and all of your domain controllers are virtualized, none of your non-virtualized domain members will have access to a domain controller either. If you have a properly designed cluster, that third issue is not really any more risky than only having a single domain controller.

Another situation often considered is making the Hyper-V hosts into domain controllers. This bypasses all of the above concerns, but at a cost. A native Hyper-V installation cannot become a domain controller, so you’ll have to use a fully licensed installation of Windows Server 2008 R2, which introduces the drawbacks as outlined in the “Host Operating System Choices” section above.

A later section of this guide covers virtualized domain controller configuration.

Cluster Shared Volumes

In a non-Hyper-V failover cluster, shared resources are identical across the nodes. Exchange Server runs the same on all nodes and connects to the same mail store. SQL Server runs on the nodes and connects to the same database. Only one node at a time “owns” the applications and their attendant resources. When a failover or node switch event occurs, ownership is transferred to another node and is accessed from there.

Hyper-V is different. The role itself is mirrored, but individual VMs run independently on the hosts. Any given VM only exists on one host at any given time, and any given host has completely different VMs than the other nodes. With a standard clustered hard drive resource, this means that each VM must live all by itself on a LUN; since only one node can own a LUN, it would be impossible to move a VM from one host to another without moving all VMs on the same LUN. This is the problem that CSV is designed to address. Each node can access a CSV simultaneously, but will only manipulate the objects assigned to VMs it has ownership of. Additionally, if a node loses connectivity to the LUN, it will check with all the other nodes; if any others can reach the LUN, the ailing node will access the LUN through the working node over the CSV network. With these advances, it is possible to have multiple VMs on the same LUN and still utilize LiveMigration.

It is important to note that currently, CSVs are **not supported** for **any** purpose other than Hyper-V.

Once created, CSV nodes will be connected to each host as a folder under C:\ClusterStorage. They and their files can be accessed at the command prompt and through Windows Explorer. This ease of browsing and manipulation should not be used lightly.

Host Configuration

Worksheet and checklist pages are provided at the end of this document.

Pre-installation Hardware Configuration

1. Ensure that all virtualization options are enabled in the BIOS, including DMA virtualization.
2. Assign iDRAC options
 - a. NIC selection: Dedicated
 - b. Register iDRAC 6 name: On
 - c. iDRAC6 name: r<<physicalhostname>> (i.e.: rhyperv1)
 - d. Domain name from DHCP: Off
 - e. Domain name: <<AD domain name>>
 - f. Host name string: FQDN of the server (c. plus e.)
 - g. Manually set all IP info
 - h. Establish a LAN user name and password. If no password is set, the iDRAC 6 will not respond to network communications.
3. Use the RAID configurator tool (CTRL+R when prompted) to “Quick Init” and, if desired, give a name to the array.

Operating System Installation and Initial Configuration

If you will be using a Server Core installation, you may wish to refer to the Server Core section later in this document. The steps here are specifically for Hyper-V R2, although they could be followed almost without alteration for a Server Core install.

1. Insert the Hyper-V R2 SP1 DVD and boot the server to it. Unless you are installing to a hard drive with an OS already installed, there are no options to change. It is advisable to delete any pre-existing partitions entirely.
2. When the server is first booted after installation, it will come up to a command prompt in one screen and, in a native Hyper-V installation, a “Server Configuration” menu in another.
3. Set the server’s time and time zone using #9 in Server Configuration or at the command-line:

```
> TIMEDATE .CPL
```

4. Change the name of the server and reboot. This can be done using option #2 in Server Configuration or at the command-line:

```
> NETDOM RENAMECOMPUTER %COMPUTERNAME% /NEWNAME:<<newname>>  
> SHUTDOWN /R /T 0
```

Network Design

Microsoft Failover Clustering requires that, on a single physical server, all NICs that will participate in cluster communications must be assigned to unique subnets. If more than one NIC exists in the same subnet and is allowed to participate in cluster communications, only one will be recognized by the Failover Clustering system.

The optimal NIC assignment is to use one dedicated subnet for each of these purposes:

- Management of the cluster and physical servers (10.1.4.0/24 in this document)
- Cluster communications, such as heartbeat information (10.0.0.0/24 in this document)
- Cluster Shared Volumes traffic (10.0.2.0/24 in this document)
- LiveMigration traffic (10.0.1.0/24 in this document)
- iSCSI traffic (10.1.2.0/24 in this document)

Broadcom NICs with current drivers can be placed in a special iSCSI HBA mode, so if you will be utilizing hosts with mixed NIC manufacturers, ensure you plan for the Broadcom cards to be utilized for iSCSI.

Initial Network Configuration

This section is primarily intended to get the server onto the network so it can be accessed by remote tools and so it can access networked resources. You may follow these steps to almost completely configure the system's networking, but it may be easier to only configure the management NIC at first and wait for the simpler tools that will be installed later to finalize the networking layout.

1. On each physical switchport that will carry iSCSI traffic: enable flow control, disable unicast storm control, and either disable STP or enable PortFast. At the switch level, enable Jumbo Frames (this will require a restart of the switch).
2. Determine the mapping of each NIC's logical name to the physical adapter it represents:
 - a. Plug in only one network cable
 - b. Run **IPCONFIG /ALL**. The connected adapter will show an autoconfigured address while all others will show "Media disconnected". Make note of the mapping.
3. Rename the management adapter (in this document, the top left PCI card is used for this purpose)

```
> NETSH INTERFACE SET INTERFACE "Local Area Connection  
<<number>>" NEWNAME="PCI Top Left (Management)"
```

4. This document uses the following names, including parentheticals, for easy identification:
 - "Onboard 1" – will be used as a virtual switch and will disappear from IPCONFIG
 - "Onboard 2" – will be used as a virtual switch and will disappear from IPCONFIG
 - "Onboard 3" – will be used for iSCSI and will disappear from IPCONFIG
 - "Onboard 4" – will be used for iSCSI and will disappear from IPCONFIG
 - "PCI Bottom Left (LiveMigration)"
 - "PCI Top Right (Cluster)"
 - "PCI Top Left (Management)"
 - "PCI Bottom Right (CSV)"

5. Configure IP addresses. At minimum, configure for the management adapter. Others can wait for simpler tools to be installed. It is recommended to not use the Server Configuration menu because it is not very tolerant of cards without gateways or DNS servers. Broadcom NICs that will be converted to iSCSI HBA and any NICs that will be converted into Hyper-V virtual switches will lose any IP and naming information entered with NETSH commands and will completely disappear from IPCONFIG. It may be advantageous to give them a minimal configuration at this time to avoid confusion later and because a NIC with a static IP address will not delay server boots while trying to obtain DHCP information.

- Setting IP information for adapters that will need traffic routed:

```
> NETSH INTERFACE IP SET ADDRESS "PCI Top Left (Management)"  
STATIC 10.1.4.25 255.255.255.0 10.1.4.1 1
```

ORDER OF ITEMS AFTER THE "STATIC" KEYWORD IS: IP, SUBNET MASK, GATEWAY ADDRESS, METRIC. METRIC IS MOSTLY IGNORED SINCE THERE WON'T BE MULTIPLE ADAPTERS IN A SINGLE SUBNET, BUT IT IS STILL REQUIRED.

- Setting IP information for NICs with unrouted traffic (CSV, cluster, iSCSI, LiveMigration):

```
> NETSH INTERFACE IP SET ADDRESS "Onboard 3" STATIC 10.1.2.25  
255.255.255.0
```

6. If you wish to add a DNS server to an adapter that should not register with a DNS server (not needed, recommended, or very useful, but included for completeness' sake):

```
> NETSH INTERFACE IP SET DNSSERVERS "Onboard 3" STATIC 10.1.0.25  
NONE
```

IF YOU WISH TO ADD MORE DNS SERVERS TO THIS ADAPTER, FOLLOW THE NEXT STEP IDENTICALLY. THE PURPOSE OF RUNNING THE ABOVE COMMAND WAS TO OVERRIDE THE DEFAULT BEHAVIOR OF REGISTERING WITH A DNS SERVER BY ADDING THE "NONE" KEYWORD. ONCE THAT IS SET, ADDING ADDITIONAL DNS SERVERS WILL NOT CHANGE IT.

7. Set DNS server information for NICs that should register, which for purposes of this document is only the management NIC. No other adapters have any need for a DNS server.

```
> NETSH INTERFACE IP ADD DNSSERVERS "PCI Top Left (Management)"  
10.1.0.25
```

- Repeat this step until all DNS servers are added

8. Adjust the MTU for all NICs connected to a switch with jumbo packets enabled. This could be skipped for Broadcom cards that will be converted to HBA mode, but is useful to set for testing.

```
> NETSH INTERFACE IP SET SUBINTERFACE "Onboard 3" MTU=9000  
STORE=PERSISTENT
```

THE MATH FOR PACKET SIZE IS INCONSISTENT. A JUMBO PACKET IS ACTUALLY 9014 BYTES. THE ABOVE IS THE PROPER SETTING, BUT THE FOLLOWING PING TEST USES A LOWER NUMBER BECAUSE A 9000 BYTE PING WILL FAIL.

- Test: `PING 10.1.2.200 -L 8000 -F`. If there is no mention of fragmentation, then the configuration change was accepted. Otherwise, check that jumbo frames is enabled on your switch and that it was rebooted after the change. Ensure that the ping target is in the same subnet as the adapter you changed the MTU on. Ensure that you used “SUBINTERFACE” and not “INTERFACE” when changing the MTU.
9. Disable NetBIOS over TCP/IP for all adapters except the management NIC and cards that will be converted to Hyper-V virtual switches. This can also be skipped for Broadcom cards that will be converted to iSCSI HBA mode.

```
> WMIC NICCONFIG GET IPAddress, index, TcpipNetbiosOptions
> WMIC NICCONFIG WHERE INDEX=<<#>> CALL SetTcpipNetbios 2
```

- You will need to re-enter the second command for each NIC to be modified, changing the “<<#>>” to match the indexes as indicated by the first command.
 - A return code of 100 means the command did not work. This usually only occurs if you attempt to set the option back to its default of 0. Open REGEDIT and navigate to `HKLM\System\CurrentControlSet\Services\NetBT\Parameters\Interfaces\Tcpip_<<GUID>>`. You can use the cues in these registry keys to determine which NIC is which, you can reset the `TcpipNetbiosOptions` to 0 for all NICs and then start this step over, you can use the second command above to assign a 1 to all the NICs you want to reset so you can locate them in the registry more easily, or you can download the `NVSPBIND.EXE` tool and run it with the `-N` switch to determine the adapter-to-GUID mapping.
10. Enforce the lowest metric on the management adapter:

```
> NETSH INTERFACE IP SET INTERFACE="PCI Top Left (Management)"
METRIC=1
```

11. Join the domain and reboot using #1 in Server Configuration or at the command-line:

```
> NETDOM JOIN %COMPUTERNAME% /DOMAIN:<<yourdomain\adminuser>>
/PD:<<adminpassword>>
> SHUTDOWN /R /T 0
```

IF YOU TYPE ONLY AN ASTERISK (*) AFTER /PD: INSTEAD OF A PASSWORD, YOU WILL BE PROMPTED TO ENTER THE PASSWORD AND IT WILL NOT BE ECHOED TO THE SCREEN AS YOU TYPE

12. With the server on the network and joined to the domain, now would be an optimal time to copy and install any updated drivers, BIOS updates, etc. For Broadcom NICs, install all available features. For Intel NICs, the SNMP features will not be installable unless you first install the Windows SNMP feature. SNMP is not required for Hyper-V, but if you intend to set up any SNMP traps or if you’re going to install Dell’s OpenManage or something similar that benefits from SNMP, this may be desirable. To install Windows SNMP:

```
> START /W OCSETUP SNMP-SC
```

Enable Remote Access and Install Prerequisite Components

1. Enable administrative remote desktop with #7 in Server Configuration or at the command-line:

```
> CSCRYPT C:\WINDOWS\SYSTEM32\SCREGEDIT.WSF /AR 0
```

2. Enable management from remote MMC (note that if you install Core Configurator, it may overwrite the firewall settings when its Firewall section is first opened). Use #4 in Server Configuration (run through all lines; a reboot will be required partway through) or at the command-line:

```
> NETSH ADVFIREWALL FIREWALL SET RULE GROUP="remote
administration" NEW ENABLE=YES
> DISM /ONLINE /ENABLE-FEATURE /FEATURENAME:NetFx2-ServerCore
/FEATURENAME:MicrosoftWindowsPowerShell
/FEATURENAME:ServerManager-PSH-Cmdlets
/FEATURENAME:BestPractices-PSH-Cmdlets
> SHUTDOWN /R /T 0
```

DESPITE THE FEATURE NAME, YOU CANNOT REMOTELY USE SERVER MANAGER. YOU CAN USE COMPMGMT.MSC.

After the reboot:

```
> POWERSHELL
  o Set-ExecutionPolicy -ExecutionPolicy RemoteSigned
  o Configure-SMRemoting.ps1 -FORCE -ENABLE
  o EXIT
```

FEATURE NAMES ARE CASE-SENSITIVE. IF YOU RECEIVE ERRORS, CHECK SPELLING AND HYPHENATION.

3. If you will be using Dell OpenManage or another tool that uses WinRM, run the following. For full usage of OpenManage, you will need to set up a server certificate and enable SSL on WinRM, which is beyond the scope of this document. This step is not necessary if you went through all of the options in Server Configuration menu item #4.

```
> WINRM QUICKCONFIG
> SC CONFIG WINRM TYPE= OWN
```

THE SPACE BETWEEN "TYPE=" AND "OWN" IS REQUIRED

4. To enable read-only access to Device Manager, perform the following steps from a computer with a full Windows GUI:
 - a. Go to Start and run MMC.EXE
 - b. Go to File->Add/Remove Snap In...
 - c. Double-click "Group Policy Object Editor"
 - d. Click Browse and find the name of the new server
 - e. Navigate to **Computer Configuration\Administrative Templates\System\Device Installation** and enable "Allow remote access to PNP interface"

ONLY READ-ONLY CAPABILITY IS POSSIBLE FOR REMOTE DEVICE MANAGER ACCESS.

Advanced Network Configuration

At this point, the .Net Framework has been installed, so the Broadcom utility is accessible. This utility is not accessible until you install the complete driver package available from Dell (or your server manufacturer, or even direct from Broadcom). If you are installing from a Dell OpenManage or repository DVD, ensure it is recent enough (December 2010).

1. Navigate to `C:\Program Files\Broadcom\BACS` and run `BACS.EXE`.
2. Ensure the view is set to "By Connection".
3. Select the highest level container for a Broadcom card to be used in iSCSI. You may need to use IPCONFIG to compare the physical names to ensure you are working with the correct NIC.
4. Select the "Configurations" tab. Expand "Resource Reservations". Ensure "First Come First Served" is checked. Under "Pre-Allocated Resources", uncheck everything except "iSCSI Offload Engine".
5. Click Apply. This will cause two delays: first, the UI is unresponsive while the configuration is saved. Second, the UI will be partially responsive, but then the line item beginning with a green icon representing the physical NIC will disappear. This NIC will no longer be visible in Windows tools such as IPCONFIG. It may still cause name and IP collisions; your mileage may vary. Checking the NDIS box in this section will return it to normal behavior.

WHEN A BROADCOM CARD IS CONVERTED TO iSCSI HBA MODE IN THIS FASHION, IT WILL STAY IN THAT MODE EVEN IF THE OPERATING SYSTEM IS REINSTALLED, ALTHOUGH ANY OTHER CONFIGURATION SETTINGS ARE LOST. YOU MUST BE ABLE TO ACCESS THIS TOOL IN ORDER TO RESET IT TO NORMAL FUNCTIONALITY.

6. Select the remaining subtree item for the NIC from step 5. On the "Configurations" tab, under IPv4 configuration, set IP information as desired. Change the MTU to 9000.
7. Repeat steps 3-6 for any other Broadcom adapter to be used in iSCSI.
8. Select the highest level container for a non-iSCSI Broadcom card. Switch to the "Configurations" tab. Uncheck the "iSCSI Offload Engine". Click Apply. This will cause the same delays as from step 5, but this time the subtree item with the green icon will remain while the blue item disappears.
9. Access the remaining subtree item and its "Configurations" tab. Change the MTU to 9000.
10. Repeat steps 8-9 for other non-iSCSI Broadcom cards.
11. If you have Intel NICs and they are attached to a switch with jumbo packets enabled, you can use the Intel utility to set the driver accordingly. Navigate to `C:\Program Files\Intel\DMIX\CL`.

```
> PROSetCL.EXE Adapter_Enumerate
```

Note the indexes of relevant Intel cards.

```
> PROSetCL.EXE Adapter_SetSetting <<#>> *jumbopacket "9014 Bytes"
```

THIS IS ALSO THE TOOL THAT WOULD BE USED TO ENABLE VMQ ON AN INTEL CARD. TO DATE, ENABLING VMQ ON AN INTEL NIC CAUSES WINDOWS TO HALT WITH AN IRQL_NOT_LESS_OR_EQUAL BLUE SCREEN WHEN THAT NIC IS CONVERTED TO A HYPER-V VIRTUAL SWITCH. THE TROUBLESHOOTING SECTION SHOWS HOW TO RECOVER FROM THIS, BUT IT IS DIFFICULT AND WILL NOT ALLOW YOU TO USE VMQ. DO NOT ENABLE THE VMQ OPTION.

Installing Dell OpenManage

This requires a recent DVD from Dell with the OpenManage software. Earlier editions were included on the Systems Build and Update Utility DVD. Recent editions have been separated onto their own DVD titled “Dell OpenManage Server Administrator Managed Node”. These instructions all assume that the DVD drive is assigned letter “D:”.

ANY TIME A DELL SYSTEMS BUILD AND UPDATE UTILITY DVD IS PRESENT AT BOOT, THE SERVER WILL BOOT TO IT WITHOUT PROMPTING, REGARDLESS OF HOW THE BOOT ORDER IS ESTABLISHED IN BIOS.

1. If you have not already installed SNMP, do so now:

```
> START /W OCSETUP SNMP-SC
```

2. If you have Intel cards and have not yet installed the driver:

```
> D:\SYSMGMT\srvadmin\support\intel_snmp_64bit\setup.exe
```

3. If you have Broadcom cards and have not yet installed the driver:

```
> D:\SYSMGMT\srvadmin\support\brcm_snmp_64bit\setup.exe
```

4. Run the OpenManage prerequisites check:

```
> D:\SYSMGMT\srvadmin\windows\PreReqChecks.exe /s
```

5. Copy the results from %TEMP% to a network location. From a computer with a web browser, open the “omprereq.htm” file, fix any issues, and repeat from step 4 as necessary. WinRM on HTTPS requires a certificate to be installed to the server, which is beyond the scope of this document. You may safely ignore this warning.

6. Navigate to `D:\SYSMGMT\srvadmin\windows\SystemManagement` and run:

```
> MSIEXEC /I SYSMGMT.MSI
```

IF YOU INTEND TO INSTALL AN SSL CERTIFICATE FOR WINRM, CHOOSE “CUSTOM” AND INSTALL EVERYTHING. OTHERWISE, CHOOSE “TYPICAL”.

7. **Ensure all DVDs are removed at the end of installation!**

8. Open the firewall for OpenManage:

```
> NETSH ADVFIREWALL FIREWALL ADD RULE NAME="DellOpenManage"  
DIR=IN ACTION=ALLOW PROTOCOL=TCP LOCALPORT=1311
```

Or:

```
> NETSH ADVFIREWALL FIREWALL ADD RULE NAME="DellOpenManage"  
DIR=IN ACTION=ALLOW PROGRAM="C:\Program Files  
(x86)\Dell\SysMgmt\iws\bin\win32\dsm_om_connsvc32.exe"
```

9. Test from a remote computer’s browser: <https://<<servername>>:1311>

Install Core Configurator and Finalize Network and Host Configuration

1. Acquire Core Configurator: <http://coreconfig.codeplex.com/>
2. If you downloaded the ZIP file, locate the downloaded file in Windows Explorer. Right-click on it and choose “Properties”. Unblock the file.
3. Extract the contents of the ZIP file to a location reachable by the server. “Install” it by copying the to a local directory on the server. Example:

```
> ROBOCOPY I:\Software\CoreConfigurator C:\CoreConfig
```

4. Usage: navigate to `C:\CoreConfig` and run `Start_CoreConfig.wsf`
5. Open Core Configurator. Go to “Control Panel”, “Firewall”, then “Rule Configuration”. Enable all except “File and Printer Sharing Management” and “Remote Schedule Task Management”.
6. Set up any remaining unconfigured NICs using Core Configurator or the NETSH commands from the Initial Network Configuration section.

DO NOT USE GATEWAYS ON NICs DEDICATED FOR CLUSTER COMMUNICATIONS, CSV, LIVE MIGRATION, OR iSCSI.

METRIC, MTU, AND DNS REGISTRATION (NONE/PRIMARY/BOTH OPTIONS) CANNOT BE SET WITH CORE CONFIGURATOR.

7. This is the optimal point to update BIOS, drivers, and any other non-Hyper-V applications or settings.

Equallogic Host Integration Toolkit Installation

This document only covers connecting to a SAN that has already been configured.

1. Obtain the latest 64-bit edition (3.5.1 as of this writing) of the Equallogic Host Integration Toolkit for Windows from <http://support.equallogic.com> (a support login is required).
2. Copy the installation file to a temporary location on the server and execute it.
3. Accept all defaults; answer “Yes” to the firewall, iSCSI, and MPIO questions.
4. On the last dialog, leave the box checked for remote configuration and allow the server to reboot.
5. If you have not yet created a LUN on the Equallogic to use as the cluster quorum, do so now. It can be of any size, but it is recommended that it be at least 500MB, as it must be formatted with NTFS and this is the minimum efficient size for that volume format. The quorum disk can be used for other purposes, although that is not recommended. It cannot be converted to a CSV.
6. If step 4 was followed, the Remote Setup Wizard runs automatically after the server reboots.

TO ACCESS LATER:

RUN “**ISCSICPL**” AT THE COMMAND-LINE OR ACCESS CORE CONFIGURATOR, CLICK “NETWORK SETTINGS”, THEN “iSCSI CONFIGURATION”. SWITCH TO THE DELL EQUALLOGIC TAB AND CLICK THE RSW BUTTON. YOU CAN DIRECTLY ACCESS THE UTILITY AT C:\PROGRAM FILES\EQUALLOGIC\BIN\REMOTESETUPWIZ.EXE”

7. Set the dot for “Configure MPIO settings for this computer”. Exclude all subnets except those dedicated for iSCSI traffic.
8. Run the RSW again (check the note under step 6) and choose “Configure this computer to access a PS Series SAN”.
9. Click “Add Group” and use the following information:
 - a. Group Name: <<*the name of your SAN group*>>
 - b. Group IP: <<*the IP of the SAN group, NOT a member IP*>>
 - c. CHAP credentials as configured; you will probably use the same for both sections
 - d. You probably do not wish to select “Use CHAP credentials for iSCSI discovery”.

IF YOU INTEND TO USE THE ADVANCED CHAP OPTIONS, IT IS ADVISABLE TO BEGIN WITH THE SIMPLEST SET UP FIRST AND, ONCE YOU ARE CERTAIN YOU CAN ESTABLISH CONNECTIVITY, TIGHTENING CHAP USAGE LATER.

10. Open the SERVICES.MSC console on a computer with a GUI. Right-click “Services (Local)” in the top of the left pane and choose “Connect to another computer”. Enter the name of the server and click OK. Once connected, change the EqlReqService account to use a domain account. That account must have local administrator rights, but does not need special domain rights.

IF YOU NEED TO CHANGE LOCAL ACCOUNTS ON THE SERVER, YOU CAN REMOTELY CONNECT WITH COMPMGMT.MSC.

Storage Setup for Hyper-V and Cluster Shared Volumes

The instructions in this section can be used any time you want to create new storage on your SAN for Hyper-V. It also provides the basic steps that would be followed if you were setting up a LUN for a Windows host to directly connect to, such as for SQL Server or Exchange.

If you are adding a host to an existing cluster, it is advisable to follow the “Setup Hyper-V and Clustering” section first, and then return to these steps. While it is possible to set up the iSCSI targets first, any CSVs this server attempts to connect to will cause it to generate a substantial number of VDS errors and will dramatically slow down anything related to the VDS service, such as Disk Management.

1. Use the Equallogic tools to build a LUN. It is also possible to use Microsoft’s Storage Manager for SANs, although the security setup from that tool is lacking. Ensure you always check the option for “Allow simultaneous connections from initiators with different IQN names”, or you will be unable to use the LUN for a CSV.
2. Run **ISCSICPL.EXE** or access iSCSI or access iSCSI in Core Configurator. The Targets tab should list LUNs on the array. If not, click Refresh.
 - a. If targets still don’t appear, go to the Discovery tab. Click “Discover Portal...” and verify the IP address of the Equallogic group (NEVER use member addresses, even if you only have one member). If you used any special security configurations on the Equallogic, enter them here.
3. The final characters of each item in the list are the common name of the LUN. To connect to a LUN:
 - a. Highlight a LUN and click “Properties” (for Intel cards, click Connect and ignore b, d, and f)
 - b. Click “Add Session”
 - c. Check the “Enable multi-path” box and click Advanced
 - d. In “Local Adapter”, choose a Broadcom card. Drop-down “Initiator IP” and ensure one of the iSCSI IPs is here. If not, this isn’t an iSCSI Broadcom; choose another Broadcom card.
 - e. Check “Enable CHAP log on”. Fill in authentication as set in Equallogic unit.
 - f. OK back to the Properties sheet. Repeat b-e for each iSCSI Broadcom card.
 - g. OK back to the Targets screen

EQUALLOGIC DOCUMENTATION STATES THAT STEP C. IS UNNECESSARY WHEN THE HOST INTEGRATION TOOLKIT IS INSTALLED

4. From a computer with a GUI, open Computer Management (COMPMGMT.MSC). Right-click “Computer Management (Local)” and choose “Connect to another computer”. Enter the name of the server and click OK. Expand Storage and go to Disk Management.
5. One disk will show as a zero-sized offline volume. This is the VDS/VSS control volume. Ignore it.
6. All other disks will be the data LUNs. Right-click and select “Online” for each LUN this server is to connect to that is not already part of a CSV (the option will be unavailable for those). For a new cluster installation, at least connect to the quorum LUN.
7. If the LUN has not been initialized, right-click it and select Initialize. Do not use GPT unless you know what it is and have a definite need for it.

NON-NTFS LUNS, SUCH AS VMFS, MAY SHOW AS UNINITIALIZED. INITIALIZING THEM WILL DESTROY THEM.

8. This step sets up a new partition's sector alignment to match the stripe size of an Equallogic array. In older Windows versions, the alignment was typically quite mismatched, resulting in performance degradation. At least as of 2008 R2, it tends to create them with 1024kb alignments. Because 1024 is evenly divisible by the Equallogic stripe size of 64kb, this step may no longer provide any performance benefits. However, it doesn't hurt anything either.

```
> DISKPART
> LIST DISK
```

You will be presented with a list of known disks. Identify which is the newly created LUN; refer to Disk Management from the GUI computer if necessary.

```
> SELECT DISK <<#>>
> LIST PARTITION
```

If you have selected a previously formatted disk, the "Offset" field shows the current alignment. Anything divisible by 64 should be acceptable.

```
> CREATE PARTITION PRIMARY ALIGN=64
```

THIS WILL DESTROY ANY EXISTING PARTITIONS ON THE SELECTED DISK

9. You can now format the volume using Disk Management (refreshing the view may be necessary) or at the command-line.

UNLIKE A LOCAL DISK, THE ONLY TIME THE NAME OF A VOLUME ON A SAN LUN CAN BE SET AS IT FORMAT TIME.

IT IS POSSIBLE TO OVERRIDE DEFAULT FORMAT SETTINGS AND SELECT A 64KB SECTOR SIZE WHEN FORMATTING AN NTFS VOLUME. BECAUSE THESE WILL MOSTLY HOLD LARGE .VHD FILES, THE SLACK SPACE LOSS WILL BE MINIMAL, BUT IT IS NOT KNOWN IF ALIGNING THE SECTOR SIZE TO THE STRIPE SIZE PROVIDES ANY BENEFITS.

Setup Hyper-V and Failover Clustering

Setting up Hyper-V and Failover Clustering is pretty simple. If you omit the “START /W” portion of these commands, they will run asynchronously, but there is really no benefit to this. Note that all feature names are case-sensitive.

1. The Hyper-V role is automatically enabled when you install Hyper-V natively. For Server Core:

```
> START /W OCSETUP Microsoft-Hyper-V
```

THE SERVER WILL REBOOT TWICE AFTER ENABLING THE ROLE ON SERVER CORE.

2. Install the Failover Clustering Role (not available on Windows Server Standard Edition). In Server Configuration, choose option #11, or at the command-line:

```
> START /W OCSETUP FailoverCluster-Core
```

3. If following this guide straight through, all Windows components have been installed. This would be an optimal time to use Core Configurator to connect to your WSUS server and install updates.

Setup Hyper-V Virtual Networking

These steps will create the virtual switch(es) that the virtual NICs of your VMs will “plug in to”. While Hyper-V will allow it, it is not advised to create virtual networks on NICs that will be used for any other purpose. The virtual switches are layer-2 with VLAN trunking capability.

1. Run `IPCONFIG /ALL` at the server console. Ensure that networking is set up as desired. If you have not already verified which logical adapters (i.e. “Intel® 82576 Gigabit Dual Port Network Connection #2”), now is the time.
2. Use a computer with the Hyper-V Manager component of Remote Server Administration Tools installed (access <http://www.microsoft.com/downloads> and search for Remote Server Administration Tools if you have not already acquired them). In Hyper-V Manager, use the “Connect to Server...” option to connect to the server.
3. Ensure the host you are working on is selected, and click “Virtual Network Manager...”
4. With “External Network” highlighted, click “Add”.
5. Match the logical network card you wish to use and name its network accordingly (i.e. “VMNet#”)

IF YOU WILL BE CREATING MULTIPLE VIRTUAL SWITCHES, SIMPLY SELECT THE “NEW VIRTUAL NETWORK” LINK AGAIN BEFORE HITTING APPLY OR OK. THIS CAN SAVE A SUBSTANTIAL AMOUNT OF TIME.

IT IS NOT ADVISABLE TO USE THESE NICs FOR ANYTHING ELSE, SO IT IS BEST TO CLEAR THE “ALLOW MANAGEMENT OPERATING SYSTEM TO SHARE THIS NETWORK ADAPTER” CHECKBOX. IF YOU LEAVE THIS CHECKED, THE CARD IS STILL CONVERTED TO A VIRTUAL SWITCH, BUT HYPER-V INSTALLS A VIRTUAL NIC IN THE HOST OPERATING SYSTEM AND CONNECTS IT TO THE NEWLY CREATED VIRTUAL SWITCH. IF THAT NIC NEEDS TO USE A VLAN, SELECT IT ON THIS DIALOG. THIS SETTING CAUSES PROBLEMS ON SOME NICs, NOTABLY OLDER INTEL CARDS. IF THE UI HANGS AND EVENTUALLY FAILS, IT IS BECAUSE THE NIC CANNOT HANDLE THE SETTING OR BECAUSE YOU HAVE NO OTHER PATH TO THE SERVER (LIKE A DEDICATED MANAGEMENT NIC). FOLLOW THE TROUBLESHOOTING SECTION TO RESET VIRTUAL SWITCHES.

6. After all virtual switches are created and you've OK'd out, you should no longer be able to see the physical NICs used for virtual switches by using NETSH or Core Configurator. If you can, it is most likely because you selected the option to allow the management operating system to use them. If that was intentional, you can use NETSH or Core Configurator to rename the host's virtual NICs.

Creating or Joining a Cluster

1. Ensure all nodes are set up as desired, with Windows Updates, BIOS, firmware, and drivers brought up to current. For optimal results, all physical servers should match.
2. From a computer with RSAT tools installed, run Failover Cluster Manager. If this is a new cluster, click "Validate a Configuration". Enter the names of all servers to be validated. If this test is run on an existing cluster, it will cause a brief service interruption during some of the tests that will cause guest VMs to shut down.
 - a. Expected warnings:
 - i. Multiple NICs on one subnet. This will occur if the iSCSI cards are not in HBA mode or if you have left your virtual switch NICs accessible to the host and did not move them into separate subnets. Because these networks will not be used for cluster communications, this is an acceptable warning. Any other NICs should be properly segregated.
 - ii. If you have not connected to SAN LUNs, there will be storage warnings.
 - iii. Some drivers may be unsigned, notably AMD IOMMU. Update if possible.
 - b. Saving the report is possible, although it is wise to run another validation test after the cluster has been created and save that one.
3. For a new cluster, start creation from the last page of the validation wizard or from the Create link.
 - a. The "Cluster Name" will act as a physical computer, complete with a DNS entry
 - b. Create only one Administrative Access Point, and create it on the dedicated management network (10.1.4.0/24)
4. To join a server to an existing cluster, connect to the cluster in Failover Cluster Manager. Right-click on the cluster name and choose "Add Node...", and follow the wizard. Only use the validation tool if the VMs are shut down or in saved state.
5. In Failover Cluster Manager, expand the "Networks" section. Highlight each cluster network then expand one of the nodes in the center pane. This will show the IP address of the member NIC. Based on that, rename the cluster network accordingly (using the Rename link on the right, double-clicking the network in the left pane, or highlighting it in the left pane and pressing F2).
6. If your iSCSI cards are not in HBA mode, at least one cluster network will have been created for them. Right-click on them in the left pane and choose "Properties". Set the radio button to "Do not allow cluster network communication on this network".
7. For Cluster, CSV, and LiveMigration networks, access their property panes as in step 6 and ensure the radio is set to "Allow cluster network communication on this network" and that the checkbox for "Allow clients to connect through this network" is cleared. This should automatically be the case if you did not enter a gateway address.

8. For a new cluster, if you have already built the quorum LUN and attached all the hosts to it, it has probably already been configured automatically as the quorum volume. If you did not take these steps, created the LUN now as instructed in the Storage Setup for Hyper-V and Cluster Shared Volumes section. Once the quorum LUN is established and all servers are connected to it, access the Storage node in Failover Cluster Manager and click “Add a disk”. The first added disk will automatically become the storage volume.
9. Despite any names you’ve set, the cluster will automatically determine which networks to use for its operations. This will probably work out just fine, but if you want to ensure it behaves as you’ve dictated, then follow these steps:

```
> POWERSHELL
> Import-Module FailoverClusters
> Get-ClusterNetwork | ft Name, Metric, AutoMetric, Role
```

The “Role” column indicates what access level the networks have. 0 means the network is not used by the cluster. 1 means only cluster traffic is allowed. 3 means both cluster and client traffic are allowed. The Metric column indicates how the network will be used. For “1” networks, the automatic metric system starts at 1000 and auto-increments by 100. The object with the lowest is for CSV, the second lowest is for LiveMigration, and the highest is for cluster communications. To override the automated system, use the following commands to manually assign metrics to these networks with numbers below 1000 but in the previously described order:

```
> (Get-ClusterNetwork “CSV”).Metric=500
> (Get-ClusterNetwork “LiveMigration”).Metric=600
```

Running the first `Get-ClusterNetwork` command (arrow-up to retrieve it) should now show the networks in the intended order.

For LiveMigration, this step only sets the **preferred** network; it does not guarantee that LiveMigration will **only** use this network. For most deployments, this should be fine. The cluster will use the designated network for LiveMigration transfers unless it becomes unavailable; then it will try the others. The other networks will generally not be carrying meaningful amounts of traffic, so having the backups available is generally preferable. If you wish to override or tune this behavior, instructions are provided in a special section.

Establishing and Using Cluster Shared Volumes

1. For a new cluster or a cluster that has never used CSVs, in Failover Cluster Manager, left-click the node. In the center or right pane, click “Enable Cluster Shared Volumes...”
2. If not already done, create LUNs on the SAN and connect to them using the steps from the “Storage Setup for Hyper-V and Cluster Shared Volumes” section. Because CSVs can be accessed by any host in the cluster simultaneously, it is not necessary to size the LUNs for a single VM. Make the CSV LUN as large as you feel comfortable with – consider snapshotting, maximum NTFS restrictions, etc.
3. In Failover Cluster Manager, go to the Storage node. Click “Add a disk” and add the new disk in. You can use its Properties dialog to give it a more meaningful name. At this point, the disk is a traditional cluster resource, owned by only one node.
4. Switch to the Cluster Shared Volumes node. Click “Add storage”. Check the disk you want to use as a CSV and click OK.
5. If you wish to rename the volume, connect to the console of the cluster or any host (example):

```
> CD\ClusterStorage  
> REN Volume1 vm_lightduty
```

Final Cluster Validation

If you’ve been following this guide straight through, the cluster has been completely configured up to this point and is ready for final testing and deployment. If you ever need to contact Microsoft Product Support Services, they will require that your cluster have passed their validation tests. Because it requires some brief but impacting downtime on the cluster, this is the optimal time to run the tests.

Any configuration change to the cluster or its members, including adding a new CSV, causes earlier validation tests to become outdated, so it is recommended that you keep your validation reports up-to-date in the event of a failure. This will require VMs to be shut down and/or saved, so schedule maintenance windows for this process.

The validation is run from Failover Cluster Manager. There is a Validate This Cluster link in the center pane. You can also pre-validate a cluster as described in “Creating or Joining a Cluster” section.

IF YOU VIEW THE REPORT IN INTERNET EXPLORER AND CLICK ON ANY LINK, AND THEN TRY TO SAVE THE REPORT TO .MHT, THE RESULTANT FILE WILL BE BLANK. SAVE THE REPORT FIRST.

Networking Troubleshooting

General Problems Related to Virtual Switches

Common problems are out-of-date drivers, not allowing the host to use the only NIC that was reachable remotely, and usage of NICs that simply won't work with virtual switches. If you don't have a complete understanding of the operation of virtual switches, that can also leave you with an unintentional misconfiguration that renders the virtual switches unusable. It is imperative that you understand that first.

Upon using Hyper-V Manager to designate a physical port as a virtual switch, Windows completely reconfigures the way it treats that NIC. It is now an actual switch, albeit a virtual one. That means that the physical switch that the NIC is connected to now believes it is talking to another switch, not a NIC endpoint. In simple environments, this may not mean much. If your environment is using VLANs or has trunking restrictions, then you must configure the physical switch's port to operate in trunking mode. Also, if you selected the option to allow your host to use the same adapter, it is not set to share the physical NIC. A virtual NIC is created, assigned to the host, and "plugged in" to this virtual switch. That means that if your environment is using VLANs, you must also assign that virtual switchport a VLAN (done in the same place you created a virtual switch on the NIC, right beneath the checkbox to allow the host to use the adapter). Some NICs cannot handle this mode, notably some older Intel cards. It may be possible to set a configuration option in the NICs driver that allows it to enable VLANs; since this should be handled at a different level anyway, this may not be effective.

There are, of course, problems that go beyond a simple misconfiguration on the part of the administrator. There have been some great resources available on the Internet showing how to quickly clean-up problems with virtual switches. Unfortunately, it seems that the most common problems needed more, and many of those resources have since disappeared. Hopefully, you encounter problems with your virtual switches before deployment, and if not, hopefully your problem was caused by something you can roll back, and if even that fails, hopefully you are using a cluster so you can work on one host while you repair the other. The process included here is somewhat lengthy and will completely destroy any virtual switches on a host, but the cleanup is very complete and allows you to start over. What these steps do is describe how to completely uninstall Microsoft's virtual switch driver and configuration from your server, and then reinstall them with no virtual switches created. You will then be able to attempt to address any problems that may have contributed to the failure and recreate your virtual switches. If your situation requires you to try more surgical approaches first, search the Internet for keywords "NVSPBIND" and "virtual switches" or "Hyper-V".

To completely reset the virtual switches, drivers, and settings:

```
> NETCFG -U VMS_PP  
> SHUTDOWN /R /T 0
```

The next step is the one that reinstalls the driver, which changed between initial release and Service Pack 1. The first line is run on an RTM host, the second on a Service Pack 1 host. Because it is highly

error-prone, the presentation is different. The text is presented in white. The optimal place to press TAB for auto-completion is in light green (the three right after the underscore after the first instance of “inf”). The reddish character indicates where auto-complete will stop and you must resume typing (right before “\wvms_pp.inf”).

Hyper-V R2 RTM and Windows Server 2008 R2 RTM:

```
> NETCFG -L
C:\Windows\winsxs\amd64_wvms_pp.inf_31bf3856ad364e35_6.1.7600.1
6385_none_beda85050b13680c\wvms_pp.inf -C P -I vms_pp
```

Hyper-V R2 SP1 and Windows Server 2008 R2 SP1:

```
> NETCFG -L
C:\Windows\winsxs>amd64_wvms_pp.inf_31bf3856ad364e35_6.1.7601.1
7514_none_c10b98cd0801eba6\wvms_pp.inf -C P -I vms_pp
```

Reboot again. Ensure that any external problems are corrected then recreate your virtual switches.

Recovering from Blue Screen Loop after Enabling Intel VMQ

Intel drivers give the option to enable VMQ, but once NICs in this mode are attached to a virtual switch, the server blue screens with IRQL_NOT_LESS_OR_EQUAL. There may be a way to use other settings in the Intel driver to avoid this problem, but Intel is not forthcoming with them. The only option is to not use VMQ. These steps will help you recover from the blue screen loop.

1. Boot the host into safe mode.
2. Run **REGEDIT**.
3. Navigate to HKLM\System\CurrentControlSet\Control\Class\ {4D36E972-E325-11CE-BFC1-08002BE10318}. These entries include all the NICs on your server, including the loopback adapter and the failover cluster network. Expand each node in turn, looking for a “PROSetNdi” key; adapters with this key are Intel cards.
4. Click the numeric node for an Intel card (000#). Find the *vmq entry. Set it to zero.
5. Repeat step 4 for all adapters.
6. Reboot the server and allow it to come up normally.
7. Follow the steps in the previous section to completely uninstall and reinstall the virtual switch drivers.

System Configuration Manager Virtual Machine Manager (SCVMM)

SCVMM 2008 R2 is a software package designed to help you manage your Hyper-V hosts. If you're in a mixed environment, it can also help you manage ESX and ESXi hosts. It generally consolidates a lot of the features found in the free Hyper-V Manager and Failover Cluster Manager, but does have some functionality not found in either. The primary benefits of SCVMM:

- Conversion of physical hosts to virtual guests (P->V)
- Storage migration, which allows you to move the .VHD files for a VM from one location to another. It performs the transfer while the guest is live, but causes downtime during the final transition while it cleans up the source. This cleanup time is mostly affected by how many snapshots it is necessary to merge into the source and the destination. Unfortunately, it will merge them in both locations separately and leave the VM unreachable during this process. If there are no snapshots to merge, the downtime is only a few minutes.

DO NOT USE THE ABOVE ON A DOMAIN CONTROLLER UNLESS YOU POWER IT OFF FIRST. THE MIGRATION INVOLVES SAVING THE STATE OF THE VM IF IT IS POWERED ON, WHICH CAN CAUSE CATASTROPHIC CONFLICTS ON A DOMAIN CONTROLLER.

SCVMM has a lot of positives, but it also has some annoyances.

- Unlike Hyper-V Manager, it forces you to turn on a VM before you can connect to it.
- It does not have an option to insert the Integration Services CD image so you can install manually; you must power off the VM and tell it to run an automated install process.
- If your SCVMM service pack level isn't the same as your Hyper-V host, you could get mismatched Integration Services installations.
- SCVMM runs a scan for problems before allowing you to perform a P->V operation. However, it only catches major showstoppers. Other things, such as an OS that doesn't have the right Service Pack to allow Integration Services to install, don't even generate a warning.
- For guaranteed compatibility purposes, SCVMM defaults to creating new VMs with emulated (legacy) network adapters instead of the newer and substantially more efficient synthetic adapters. This and other hardware configuration annoyances can be worked around by building your own hardware templates.

Installing the SCVMM Server

This software requires a full (not Core) Server 2008 or Server 2008 R2 installation with 2GB of RAM. It is perfectly acceptable to install it on a virtual machine. The only issue with doing so is that LiveMigrations from SCVMM involving the SCVMM machine itself will sometimes fail. Putting the host running a VM hosting SCVMM into maintenance mode will almost always fail; manually LiveMigrating it before starting maintenance mode avoids this. If the VM were to fail, this software does nothing critical that Hyper-V Manager and/or Failover Cluster Manager cannot handle.

1. Create a domain account for VM management. Give it "Log on as a service rights" either at the domain level or at the host level.

ADJUSTING “LOG ON AS A SERVICE RIGHTS” IN THE DOMAIN’S GROUP POLICY CAN HAVE UNINTENDED CONSEQUENCES, SUCH AS REMOVING IT FROM OTHER ACCOUNTS THAT HAVE IT ON LOCAL MACHINES BUT NOT IN THE DOMAIN’S GROUP POLICY. THE DOMAIN POLICY DOES NOT ADD TO THE LIST ON LOCAL MACHINES, BUT OVERRIDES IT ENTIRELY.

2. Join the VM administration domain account to the local Administrators and Distributed COM Users group on the VMM server and the Hyper-V hosts.
3. Install the IIS role with the following options: ASP.NET and its automatically selected prerequisites, IIS 6 Metabase Compatibility and its prerequisites, IIS 6 WMI Compatibility.
4. Install MBSA from <http://go.microsoft.com/fwlink?LinkId=97952>
5. Run the installer from the SCVMM disk. Use the links to download and install the configuration analyzer. Run that from the Start Menu. Correct any issues it reports.
6. Install SCVMM from the links from step 5. To keep dependencies to a minimum:
 - a. Let it use a local copy of SQL Express rather than a network SQL Server
 - b. Place the MSSCVMMLibrary share on a separate drive or VHD than the C: drive.
 - c. Use the VMM administration account to run the VMM Server service, not LocalSystem
7. “Best Practices” suggests changing the default ports that SCVMM uses. It’s probably more convenient to exercise some control over who can access your network in the first place.

Installing an SCVMM Administration Station and Configuring the SCVMM Server

1. Install from the same DVD as the server component. The management component can be installed on the SCVMM if desired. The installation is straightforward with only preferential options.
2. After installation, start the console and direct it to connect to the machine running the server.
3. If the SCVMM server has already been configured, no further steps are necessary. The client will display the installation as configured on the server.
4. If you envision having a large virtual server farm, create a new Host Group. Otherwise, the “All Hosts” default group is sufficient. Select your Host Group of choice with a left-click.
5. Click “Add Host”. Enter the name of the cluster. SCVMM will automatically detect that it is a cluster and add all hosts.
6. If you installed SCVMM on a virtual machine, you’ll need to reboot it to clear the “Incomplete Configuration” message.
7. You may stop here; the minimal configuration of SCVMM is complete.
8. To attach shared ISOs instead of copied ISOs to virtual DVD drives:

SCVMM WON’T ALLOW YOU TO CREATE A NEW VM WITH A SHARED ISO ATTACHED. YOU HAVE TO CREATE THE VM FIRST, THEN EDIT ITS PROPERTIES AND ADD THE SHARED ISO. YOU CAN CREATE A VM WITH A COPIED ISO ATTACHED.

- a. In Active Directory Users and Computers, open the properties sheet for a Hyper-V host. On the Delegation tab, choose “Trust this computer for delegation to specific services only.” Set for any authentication and click Add.
- b. Click “Users or Computers...” and choose the computer account running SCVMM. Highlight CIFS and click OK.
- c. Repeat a. and b. for other hosts.

- d. Adjust the share and NTFS permissions on the MSCVMMLibrary share and folder to give the SCVMM admin account and the host machine accounts Read access.
- e. If you are unable to use shared ISOs after following the above steps, follow steps 2-18 in the Server Core installation section of [http://technet.microsoft.com/en-us/library/cc794756\(W.S.10\).aspx](http://technet.microsoft.com/en-us/library/cc794756(W.S.10).aspx) for each Hyper-V host.

Wrapping Up

Override LiveMigration's Automatic Network Selection

In the configuration section of the guide, it was described how to select the networks that the various functions of the Hyper-V Cluster will use. In the case of LiveMigration, all it does is set the preferred network. In the event that the designate network is not available, LiveMigration will try other cluster networks until it tries one that works. In most cases, those networks will not have much traffic on them. An exception would be in the instance that the problem causing the LiveMigration network to be offline also affects iSCSI in a way that causes CSV traffic to be routed over that network. The method of changing this behavior is not exactly intuitive, most likely because it uses Failover Cluster Manager which is not a Hyper-V specific tool.

To override or tune the networks that LiveMigration uses:

1. Open Failover Cluster Manager.
2. Open the "Services and Applications" node. Select any virtual machine resource.
3. In the center pane, right-click the virtual machine and click Properties.
4. Switch to the "Network for live migration" tab. Change the ordering and accessibility of the individual networks as desired. This setting will automatically apply to all VMs on this cluster.

Special Settings for Virtualized Domain Controllers

It is vital that domain controllers never be placed in a saved state. If they detect that they have been saved, they will refuse to participate in replication and will essentially become read-only. The only recovery methods are to restore from backup or rebuild the domain controller. It is preferable to avoid these conditions. Additionally, domain controllers do not cope well with Hyper-V's automatic time synchronization.

1. Follow steps 1-3 from "Override LiveMigration's Automatic Network Selection".
2. Choose the "Settings" tab. Change the "Cluster-controlled offline action" to "Shut down" and OK.
3. Right-click on the VM in the center pane and choose "Settings".
4. In the "Integration Services" section, uncheck "Time Synchronization".
5. In the "Automatic Stop Action" section, choose "Shut down the guest operating system". Click OK.

RemoteFX Installation

This is not supported by the reference hardware in this guide. It is included for the sake of completeness. RemoteFX cannot be installed on a Server Core host. It can only be installed on a native Hyper-V host or a full installation of Windows Server.

```
> POWERSHELL
> IMPORT-MODULE ServerManager
> Add-WindowsFeature RDS-Virtualization
> Add-WindowsFeature RDS-RemoteFX
```


Using Server Core Instead of Hyper-V

In addition to installing Hyper-V directly, it is possible to use a core or full installation of Windows Server 2008 R2. This adds the ability to use features not included in Hyper-V and may be an optimal path for organizations whose Hyper-V needs will not demand all that the host hardware can offer. Installing within a full GUI installation of Windows is very simple and will not be covered. These instructions are intended to be run on a Server Core installation, but could just as easily be done from an elevated command prompt in a GUI install.

If the server is booted to the Dell Systems Build and Update Utility DVD, it can automate the installation of Server 2008 R2. This will preinstall some of the drivers, although it does not install everything that will be needed, such as the BACS utility for Broadcom cards. This utility really doesn't save as much time as it did on previous versions of Windows, so its value is questionable.

ANY TIME A DELL SERVER IS BOOTED WITH AN SBUU DISC INSERTED, IT WILL BOOT TO IT WITHOUT PROMPTING

Installation of Server Core is very simple. The only addition over a Hyper-V install is that you must select the desired operating system edition, and of course, you must have a license for it. Remember that Standard Edition cannot join a failover cluster.

Configure a Core install by following the Hyper-V instructions from the beginning of the guide. The "Server Configuration" menu is not available, so command-line alternatives have been provided.

Windows License Activation for Server Core

Do this last, after you've had sufficient time to test your installation and are certain you will be making no drastic alterations; the activation process can be painful and is not something you want to be repeating in the event of an operating reinstall or some other dramatic change.

The two activation options are KMS and MAK. KMS requires periodic, although fairly rare, connections to the KMS host, which will in turn periodically contact Microsoft's servers. The other method, MAK, does not require a connection to Microsoft's servers and is permanent until the operating system is reformatted or undergoes a significant hardware change.

KMS is the easiest; if you already have a KMS host, then shortly after being joined to the domain, your server will query for it in DNS and automatically register itself (provided that your KMS host has the appropriate level; it needs a C key to activate Datacenter, a B or C key to activate Enterprise, and an A, B, or C key to activate Standard). If you intend to make your Hyper-V host the KMS host, then follow step 1 of either the command-line or Core Configurator instructions below using your KMS key and stop. It will not go into effect until 5 hosts have contacted it for activation.

MAK is the most difficult but also the most permanent and least invasive. Activation is handled at the Microsoft licensing center and you will not need a dedicated on-premises host for it.

1. Enter your product key.

Core Configurator: Click on the "Licence Settings" button, switch to the "Licence configuration" tab

and check “Install product key”.

Command-line:

```
> SLMGR /IPK <<your-product-key-dashes-optional>>
```

2. If desired, attempt online activation. Offline activation instructions start in the next step.

Core Configurator:

Switch to the “Licence configuration tab”, check “Activate Server Core (Web) and click OK.

Command-line:

```
> SLMGR /ATO
```

3. If the above fails or you wish to perform an offline activation, continue.

Core Configurator: Switch to the “Licence information” tab and set the radio to “Display Installation ID for offline activation”

Command-line:

```
> SLMGR /DTI
```

4. You will be presented with a dialog box that contains your installation ID. It is highly recommended that you copy this key into a text editor (if you connected by RDP, you can use CTRL-C to copy the contents of the dialog to the clipboard) or write it down so that you can break the key into 6-digit groups.
5. From a computer with a GUI, run **SLUI 4** to obtain the phone number for activation.
6. When you call, you’ll receive an automated attendant. She will ask for the number from step 2 in 6-digit blocks. The most important thing to remember is:

DO NOT GET IMPATIENT WITH THE AUTOMATED ATTENDANT. LET HER FINISH EVERY WORD OF EVERY SENTENCE BEFORE YOU SAY ANYTHING. IF YOU SPEAK OR DIAL OVER HER, YOU HAVE TO WAIT FOR HER TO COMPLAIN THAT SHE DIDN’T UNDERSTAND YOU, AND THEN YOU’LL HAVE TO START THE ENTIRE BLOCK OVER.

7. Once she receives the entire installation ID, she will begin to recite another number back to you. You will enter this as follows:

Core Configurator:

On the “Licence configuration” tab, check the box for “Activate Server Core with user-provided CID” and enter the number into the single line right below it and press Apply.

Command-line:

```
> SLMGR /ATP <<extremelyLongnumberallruntogether>
```

Installation Worksheets and Checklists

Checklist and worksheet items are grouped for logical organization, not necessarily by process order.

Pre-Installation Planning

Host Name	
Cluster Name	
Management network (x.x.x.x/x)	
Management IP for this Host	
iSCSI Network (x.x.x.x/x)	
iSCSI IP(s) for this host	
Cluster Network (x.x.x.x/x)	
Cluster IP for this host	
CSV Network (x.x.x.x/x)	
CSV IP for this host	
LiveMigration network (x.x.x.x/x)	
LiveMigration IP for this host	
iDRAC Name	
iDRAC IP	

Pre-Installation Hardware Configuration and Initial Operating System Configuration

- Virtualization options enabled in BIOS
- Assign iDRAC options
 - NIC selection: Dedicated
 - Register iDRAC 6 name: On
 - Enter iDRAC6 name
 - Domain name from DHCP: Off
 - Domain name: <<AD domain name>>
 - Host name string: FQDN of the server (c. plus e.)
 - Manually set all IP info
 - Establish a LAN user name and password. If no password is set, the iDRAC 6 will not respond to network communications.
- Use the RAID configurator tool (CTRL+R when prompted) to “Quick Init” and, if desired, give a name to the array.
- Install operating system
 - Set time/date
 - Rename computer

Operating System Configuration

- Join domain
- Install SNMP
- Enable administrative remote desktop
- Enable remote MMC management
- Install PowerShell and NetFx2-ServerCore
- Set PowerShell execution policy to RemoteSigned
- Read-only access to Device Manager enabled
- Additional drivers installed
- Windows Update run
- Activated (Windows Server only)

Software Installation

- Core Configurator installed
- Equallogic Host Integration Toolkit installed
 - MPIO settings configured
 - Connected to SAN
- SCVMM installed and configured

Hyper-V and Failover Clustering

- Hyper-V role enabled (Windows Server only)
- Failover Clustering enabled
- Virtual switches setup
- Cluster created/joined
- CSVs created/connected
- Cluster validation report created and saved

Hyper-V and Server Core Command-Line Cheatsheet

Standard Command-Line

Command/Action	Result
[TAB]	Auto-complete path/filenames and some PowerShell commands
[CTRL]-[ALT]-[DEL]	Access to "Lock this computer", "Log off", "Change a password..." and "Start Task Manager"
[CTRL]-[ALT]-[END]	Access to CTRL-ALT-DEL menu via Remote Desktop Connection or Virtual Machine Connection
"Start Task Manager"->File->Run	Access to a Run prompt if the command-line interface has been accidentally closed
CMD.EXE	Starts a new instance of the command-line (useful if it has been inadvertently closed)
LOGOFF	Logs the active user off
C:\Windows\System32\sconfig.cmd	Opens "Server Configuration" (native Hyper-V)
SHUTDOWN /S /R /T #	If /S is specified, the server shuts down. If /R is specified, the server is rebooted. The number after /T indicates how many seconds to wait before shutting down; 0 for Now.
DISM.EXE /Online /?	Shows options for listing, enabling, and disabling Windows features. Features are case-sensitive
ISCSICPL.EXE	Opens the iSCSI control panel applet
POWERSHELL.EXE	Invokes PowerShell, if installed

PowerShell

A fantastic PowerShell Module for Hyper-V is available from <http://www.codeplex.com/PsHyperV>.

Command	Result
Import-Module <<modulename>>	Grants access to commands in PowerShell libraries
Import-Module ServerManager	Required for access to many other commands, such as the ?-WindowsFeature entries
Get-WindowsFeature	Shows installed Windows roles and features; similar to DISM but shows different features
Add-WindowsFeature <<featurename>>	Adds a Windows feature, case-sensitive
Remove-WindowsFeature <<featurename>>	Removes Windows feature, case-sensitive