



Microsoft® System Center

Deploying Hyper-V with Software- Defined Storage & Networking

Microsoft TechNet and the Cloud Platform Team
Mitch Tulloch, Series Editor

PUBLISHED BY
Microsoft Press
A division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2015 by Microsoft Corporation All rights reserved.

No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2014959113
ISBN: 978-0-7356-9567-2

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://aka.ms/tellpress>.

This book is provided "as-is" and expresses the author's views and opinions. The views, opinions and information expressed in this book, including URL and other Internet website references, may change without notice.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Acquisitions Editor: Karen Szall
Developmental Editor: Karen Szall
Editorial Production: Megan Smith-Creed
Copyeditor: Megan Smith-Creed
Cover: Twist Creative • Seattle

Contents

	Introduction	ix
Chapter 1	Design and planning	1
	Choosing a configuration.....	1
	Recommended hardware.....	1
	Minimal hardware.....	3
	Architecture.....	4
	Management.....	6
	Logical networking.....	11
	Storage.....	14
	Compute.....	21
	Network virtualization.....	23
Chapter 2	Deploying the management cluster	35
	Overview of the management cluster.....	35
	Virtual Machine Manager.....	37
	Service provider vs. enterprise.....	38
	Configuration walkthrough.....	38
	Procedure 1: Rack and connect management cluster hosts.....	38
	Procedure 2: Install Windows Server on MGMT01 and MGMT02.....	39
	Procedure 3: Configure TenantNetwork teams.....	40
	Procedure 4: Enable the Hyper-V role on MGMT01 and MGMT02.....	43
	Procedure 5: Configure Hyper-V virtual switches.....	43
	Procedure 6: Create folder structure for VMs and software.....	45

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

<http://aka.ms/tellpress>

Procedure 7: Configure Hyper-V settings.....	46
Procedure 8: Create management VMs.....	47
Procedure 9: Configure Active Directory Domain Services on DC01.....	49
Procedure 10: Join management hosts to the Contoso domain.....	50
Procedure 11: Configure DC02 as a secondary DC.....	51
Procedure 12: Configure remaining management VMs.....	52
Procedure 13: Install and configure WSUS on the WSUS VM.....	54
Procedure 14: Configure WSUS GPO and auto-approvals.....	56
Procedure 15: Install Windows Deployment Services.....	57
Procedure 16: Create administrative service accounts.....	58
Procedure 17: Add data and log disks to SQL01 and SQL02.....	59
Procedure 18: Add a second virtual network adapter to SQL01 and SQL02.....	60
Procedure 19: Create the guest cluster within SQL01 and SQL02.....	61
Procedure 20: Install SQL Server on SQL01 and SQL02.....	63
Procedure 21: Enable AlwaysOn high availability for SQL01 and SQL02.....	65
Procedure 22: Add a second virtual network adapter to VMM01 and VMM02.....	65
Procedure 23: Create the guest cluster within VMM01 and VMM02.....	66
Procedure 24: Create an Active Directory Domain Services container for distributed key management.....	68
Procedure 25: Install Virtual Machine Manager management server prerequisites.....	69
Procedure 26: Install Virtual Machine Manager management server.....	70
Procedure 27: Create an AlwaysOn Availability Group for VirtualManagerDB.....	72
Procedure 28: Finalize Virtual Machine Manager installation.....	73

Chapter 3	Configuring network infrastructure	75
	Configuration walkthrough	75
	Logical networks.....	76
	Network sites	77
	Procedure 1: Create a logical network and site for tenant traffic	78
	Static IP address pools	80
	Procedure 2: Create an IP address pool for the Tenant_LN logical network	81
	Procedure 3: Create a second logical network and site for the datacenter network	83
	Procedure 4: Create static IP address pools for the Datacenter_LN logical network.....	85
	Logical switches	87
	Procedure 5: Create an uplink port profile.....	90
	Procedure 6: Create a virtual network adapter port profile	92
	Procedure 7: Create a port classification	94
	Procedure 8: Create the logical switch.....	94
Chapter 4	Configuring storage infrastructure	99
	Scale-Out File Server.....	100
	Configuration walkthrough	101
	Procedure 1: Rack and connect SOFS nodes	102
	Procedure 2: Configure Baseboard Management Controllers	103
	Procedure 3: Rack and connect JBODs to SOFS nodes.....	104
	Procedure 4: Configure a BMC administrator in System Center Virtual Machine Manager	105
	Procedure 5: Create a run-as account for SetupAdmin	105
	Procedure 6: Obtain a virtual hard disk for server deployment.....	106
	Procedure 7: Create a physical computer profile.....	107
	Procedure 8: Add a PXE server to System Center Virtual Machine Manager	109
	Procedure 9: Pre-provision Active Directory accounts	111
	Procedure 10: Configure the WDS server with DHCP.....	112

Procedure 10: Download and import the Windows Server Gateway service templates into System Center Virtual Machine Manager	184
Procedure 11: Customize the Windows Server Gateway service template.....	188
Procedure 12: Configure deployment of the Windows Server Gateway service	189
Procedure 13: Configure Windows Server Gateway 001 integration with System Center Virtual Machine Manager	193
Procedure 14: Create a VM network and virtual subnets.....	197
Procedure 15: Create accounting and HR VMs	198
Procedure 16: Test VM communication	200
Procedure 17: Review network virtualization configuration	202
Procedure 18: Configure Windows Server Gateway 002 integration with System Center Virtual Machine Manager	206
Procedure 19: Create VM networks and virtual subnets.....	208
Procedure 20: Create testing and development VMs.....	210
Procedure 21: Test VM communication	211
Procedure 22: Configure NAT rules for VM networks	212
Procedure 23: Review network virtualization configuration	214

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

<http://aka.ms/tellpress>

Introduction

When you're looking at testing a new IT solution—such as implementing a software-defined datacenter that includes virtualization, networking, and storage—the best starting point is always to get advice from someone who has already done it. You can learn from experience what to do and what to avoid. That's the idea behind this book. We've gone through the work of deploying Windows Server, Microsoft System Center, and the innovations that Microsoft Azure has brought to these technologies. Our goal is to give you the step-by-step benefit of our proof-of-concept implementation to save you time and effort. And we want to show you how you can take advantage of innovation across the datacenter and the cloud to simplify your infrastructure and speed delivery of services to the business.

Transforming the datacenter

You know that IT infrastructure matters. With the right platform, you can reduce costs, respond more quickly to business needs, and take on the challenges of big data and mobility.

IT today is under more pressure than ever before to deliver resources faster, support new business initiatives, and keep pace with the competition. To handle these demands, you need a flexible, resilient infrastructure that is easy to manage and easy to scale. This means you need to be able to take everything you know and own today and transform those resources into a software-defined datacenter that is capable of handling changing needs and unexpected opportunities.

With Windows Server, Microsoft System Center, and Microsoft Azure, you can transform your datacenter. Virtualization has enabled a new generation of more efficient and more highly available datacenters for your most demanding workloads. Microsoft virtualization solutions go beyond basic virtualization capabilities, such as consolidating server hardware, and let you create a comprehensive software-defined compute engine for private and hybrid cloud environments. This flexibility helps your organization achieve considerable cost savings and operational efficiencies with a platform on which you can run the most demanding, scalable, and mission-critical of workloads.

You can find a large part of those savings and some of the best options for simplifying the datacenter in the area of storage. Microsoft's software-defined storage (SDS) capabilities enable you to deploy low-cost, commodity hardware in a flexible, high-performance, resilient configuration that integrates well with your existing resources.

Another area of savings and optimization is in networking innovation. With software-defined networking (SDN), you can use the power of software to transform your network into a pooled, automated resource that can seamlessly extend across cloud boundaries. This allows

optimal utilization of your existing physical network infrastructure, as well as agility and flexibility resulting from centralized control, and business-critical workload optimization from deployment of innovative network services. Virtual networks provide multitenant isolation while running on a shared physical network, ultimately allowing you to manage resources more effectively, without the complexity associated with managing traditional networking technologies such as Virtual Local Area Networks (VLANs).

System Center provides the unified management capabilities to manage all of this virtualized infrastructure as a whole. This software-defined model lets you pool resources and balance demand across all the different areas of the business, moving resources to the places where you need them most, increasing agility and the overall value of IT to the business.

Although the benefits of a software-defined datacenter are clear, designing and implementing a solution that delivers the promised benefits can be both complex and challenging. As with all new advances in technology, experienced architects, consultants, and fabric administrators often find it difficult to understand the components and concepts that make up a software-defined datacenter solution. We wrote this book to help.

Who should read this book?

You only have to perform a quick web search on “deploying Hyper-V,” “configuring Storage Spaces,” or “understanding Hyper-V Network Virtualization,” to realize that a wealth of information is available across Microsoft TechNet, blogs, whitepapers, and a variety of other sources. The challenge is that much of that information is piecemeal. You’ll find an excellent blog post on configuring Storage Spaces, but the networking configuration used is vastly different from the whitepaper you’ve found that guides you through configuring network virtualization. Neither of these sources align with a bare-metal Hyper-V deployment article you’ve been reading. The point here is that it’s difficult to find a single end-to-end resource that walks you through the deployment of the foundation of the Microsoft software-defined datacenter solution, comprising software-defined compute, storage, and networking, from the racking of bare-metal servers, through to the streamlined deployment of virtual machines (VMs). This book does just that.

Providing a POC deployment, this book gives the what, why, and the how of deploying the foundation of a software-defined datacenter based on Windows Server 2012 R2 and System Center 2012 R2. If you’re an IT professional, an infrastructure consultant, a cloud architect, or an IT administrator, and you’re interested in understanding the Microsoft software-defined datacenter architecture, the key building blocks that make up the solution, the design considerations and key best practices, this book will certainly help you. By focusing on a POC scale, you can implement a solution that starts small, is manageable, and is easy to control yet helps you learn and understand why we chose to deploy in a certain way and how all of the different pieces come together to form the final solution.

What topics are included in this book?

This book, or proof-of-concept (POC) guide, will cover a variety of aspects that make up the foundation of the software-defined datacenter: virtualization, storage, and networking. By the end, you should have a fully operational, small-scale configuration that will enable you to proceed with evaluation of your own key workloads, experiment with additional features and capabilities, and continue to build your knowledge.

The book won't, however, cover all aspects of this software-defined datacenter foundation. The book won't, for instance, explain how to configure and implement Hyper-V Replica, enable and configure Storage Quality of Service (QoS), or discuss Automatic Virtual Machine Activation. Yet these are all examples of capabilities that this POC configuration would enable you to evaluate with ease.

- **Chapter 1: Design and planning** This chapter focuses on the overall design of the POC configuration. It discusses each layer of the solution, key features and functionality within each layer, and the reasons why we have chosen to deploy this particular design for the POC.
- **Chapter 2: Deploying the management cluster** This chapter focuses on configuring the core management backbone of the POC configuration. You'll deploy directory, update, and deployment services, along with resilient database and VM management infrastructure. This lays the groundwork for streamlined deployment of the compute, storage, and network infrastructure in later chapters.
- **Chapter 3: Configuring network infrastructure** With the management backbone configured, you will spend time in System Center Virtual Machine Manager, building the physical network topology that was defined in Chapter 2. This involves configuring logical networks, uplink port profiles, port classifications, and network adaptor port profiles, and culminates in the creation of a logical switch.
- **Chapter 4: Configuring storage infrastructure** This chapter focuses on deploying the software-defined storage layer of the POC. You'll use System Center Virtual Machine Manager to transform a pair of bare-metal servers, with accompanying just a bunch of disks (JBOD) enclosures, into a resilient, high-performance Scale-Out File Server (SOFS) backed by tiered storage spaces.
- **Chapter 5: Configuring compute infrastructure** With the storage layer constructed and deployed, this chapter focuses on deploying the compute layer that will ultimately host workloads that will be deployed in Chapter 6. You'll use the same bare-metal deployment capabilities covered in Chapter 4 to deploy several Hyper-V hosts and then optimize these hosts to get them ready for accepting virtualized workloads.

- **Chapter 6: Configuring network virtualization** In Chapter 3, you will have designed and deployed the underlying logical network infrastructure and, in doing so, laid the groundwork for deploying network virtualization. In this chapter, you'll use System Center Virtual Machine Manager to design, construct, and deploy VM networks to suit a number of different enterprise scenarios.

By the end of Chapter 6, you will have a fully functioning foundation for a software-defined datacenter consisting of software-defined compute with Hyper-V, software-defined storage, and software-defined networking.

This book is focused on the steps to implement the POC configuration on your own hardware. Where applicable, we have included detail on design considerations and best practices and extra detail on certain features and capabilities. These are intended to ensure that you come away from this book with a rounded view of the what, why, and how when it comes to deploying the foundation of a software-defined datacenter based on Windows Server 2012 R2 and System Center 2012 R2.

Acknowledgments

The authors would like to thank Jason Gerend, Jose Barreto, Matt Garson, and Greg Cusanza from Microsoft for providing valuable guidance and contributions for the content of this book. Without their expertise and guidance, this book would not be as thorough, detailed, and accurate. Our sincere thanks go to them for their time and efforts in making this happen.

The authors would also like to thank Karen Forster for proofing and copyediting their manuscript, Deepti Dani for her work on formatting and final layout, and Masood Ali-Husein for his reviewing work on this project.

Free ebooks from Microsoft Press

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

<http://aka.ms/mspressfree>

Check back often to see what is new!

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

<http://aka.ms/HyperV1>

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *<http://support.microsoft.com>*.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: *<http://twitter.com/MicrosoftPress>*.

This page intentionally left blank

Design and planning

When it comes to implementing new technologies, especially those that will ultimately form the foundation of your datacenter platform, the design and planning phase is arguably the most important. With a poor design, even the greatest implementation skills can't prevent that deployment from failing to reach the highest levels of efficiency, reliability, and performance.

Proof of concepts (POCs) and pilot deployments, although significantly smaller than a production implementation, should still be designed, planned, and deployed in a way that reflects how a real-world environment would look. That's not to say that you couldn't modify certain elements to streamline the POC process in this non-production configuration, but you should adhere to best practices.

Choosing a configuration

During the design of this POC configuration, many questions arose. For example, how many management nodes and virtual machines (VMs) are needed? How many storage nodes and just a bunch of disk (JBOD) enclosures are needed? How many hard disks versus solid state disks (SSDs) are needed? What capacities are required of the respective disks? How many networks, ports, and switches are needed? Are more than two Hyper-V compute nodes necessary? These are just a selection of the large number of questions that we worked through when choosing the configuration for this book.

The challenge is, however, that no matter what configuration of hardware and software we chose, it's unlikely that anyone reading this will have an exact match, server for server, disk for disk, port for port, with the hardware configuration we have. That said, the book is written in such a way that even if you don't have exactly the same hardware configuration, you should still be able to confidently follow along and implement the POC on your hardware.

Recommended hardware

For this POC configuration, with the aim of keeping the complexity of building a cloud platform down to a minimum, we chose a small-scale configuration that, at a high level, consists of:

- Two modern, dual-socket x86 servers for the management nodes, each with two Remote Direct Memory Access- (RDMA-) capable 10-Gbps and four 1-Gbps network adapters along with a baseboard management controller. Management nodes also have an additional 250 GB of local storage for storing VMs.

- Two modern, dual-socket x86 servers for the storage nodes, each with two RDMA-capable 10-Gbps network adapters, along with two 6-Gbps SAS HBA ports and a baseboard management controller
- Four modern, dual-socket x86 servers for the compute nodes, each with two RDMA-capable 10-Gbps and four 1-Gbps network adapters, and a baseboard management controller (BMC)
- One modern, dual-socket x86 server for the gateway node with two RDMA-capable 10-Gbps and four 1-Gbps network adapters and a BMC
- Two 12-bay JBOD enclosures, each with four 240-GB SSDs and eight 1-TB hard disk drives (HDDs)

That's a total of nine servers and two JBOD enclosures, each with just under 9 TB of raw capacity. Each of the servers has two 10-Gbps network adapters, which you'll want to ensure are RDMA capable in case you want to evaluate the impact of RDMA on storage and network performance. RDMA-capable network adapters come in several different forms, and in this configuration, either RDMA over Converged Ethernet (RoCE) or Internet Wide-Area RDMA Protocol (iWARP) adapters would be suitable. Two 10-Gbps network adapters per server means a total of 18 10-Gbps ports and cables, respectively.

In addition to the 10-Gbps network adapters, the management, compute, and gateway nodes have four 1-Gbps network adapters, which is a total of 28 1-Gbps network ports and cables. Additionally, nine ports and cables are required to support out-of-band management with the onboard BMC. This rounds out the total at 37 cables and ports to support the 1-Gbps network.

These connections, regardless of whether they are 1 Gbps or 10 Gbps, should be distributed across multiple switches. This configuration will use two network switches to support the 1-Gbps network and two network switches to support the 10-Gbps network. Each server will have its network connections distributed across both switches, and the respective switches will be connected together to ensure redundancy.

You'll notice that we called out the use of modern dual-socket servers for the configuration. What is modern? A year old? Two? Three? Windows Server 2012 R2 was released to manufacturing in August 2013, and Windows Server 2012 was released to manufacturing in August 2012. Servers that are certified for either of these platforms (and generally, servers are certified for both) are more than adequate for this POC configuration. If you're unsure, refer to the Windows Server Catalog (<http://windowsservercatalog.com/default.aspx>) to confirm whether your hardware is certified and supported.

At the time of this writing, more than 1,500 servers are certified for both Windows Server 2012 and Windows Server 2012 R2. What if your servers are not certified? Will they not work? For a POC configuration, you should still be able to use those servers, but bear in mind that you may not be able to take advantage of some of the newer capabilities. In addition, you may not realize the same levels of performance as you would on more modern hardware that has been designed and certified for recent versions of Windows Server. That said, if you're going to

run on hardware that isn't certified on the Windows Server Catalog, ensure that your BIOS, firmware, and other hardware component-level updates have been applied.

Minimal hardware

The next question is typically whether you can build the POC configuration on smaller hardware. The answer is yes, but with some caveats. Here's the bare minimum number of servers that will allow you to follow the POC configuration and evaluate the majority of the capabilities:

- One modern, dual-socket x86 server for the management node, with two 10-Gbps and four 1-Gbps network adapters, and a BMC. The management node should have an additional 250 GB of local storage for storing VMs.
- Two modern, dual-socket x86 servers for the storage nodes, each with two 10-Gbps network adapters and two 6-Gbps SAS ports and a baseboard management controller
- Two modern, dual-socket x86 servers for the compute nodes with two 10-Gbps and four 1-Gbps network adapters and a BMC
- One modern, dual-socket x86 server for the gateway node with two 10-Gbps and four 1-Gbps network adapters and a BMC
- One 12-bay JBOD enclosure, and each JBOD with four SSDs and eight hard disk drives (HDDs)

For this minimal configuration, you'll notice that the number of management nodes is reduced to one. As you'll learn, the management node supports the VMs that will manage the overall infrastructure. These include VMs that run Active Directory Domain Services, Windows Deployment Services (WDS), Windows Server Update Services (WSUS), SQL Server instances and System Center Virtual Machine Manager servers. Many of these workloads have multiple VMs deployed to provide them with their own level of high availability. For instance, multiple SQL Server VMs allow for the deployment of a SQL Server AlwaysOn Availability Group. Placing both of these SQL Server VMs on a single management node is no longer an optimal configuration—certainly not in a production deployment. Therefore, having two management nodes with a SQL Server VM on each is preferable but not essential if you have limited hardware.

For the storage, you'll still require two physical servers for deployment of the Scale-Out File Server (SOFS). However, you could, if required, relax the requirement for two JBOD enclosures and have a single JBOD enclosure. It is still advisable to have a mix of SSD and HDD drives within the JBOD to ensure you can use storage tiering, and the ratio of SSD to HDD should still be around 1:4 (20 percent/80 percent).

In reality, you could create a single-node SOFS cluster, but it would lack redundancy. All demand would be placed on this single node instead of being evenly spread across two nodes. It is therefore a best practice to use a minimum of two nodes for your SOFS cluster for this POC configuration.

Finally, you could reduce the number of compute nodes to two. This would still allow you to create a cluster and benefit from additional resiliency while still experiencing features such as Dynamic Optimization and Live Migration. The size of your compute nodes, in terms of processor and memory, is largely determined by the workloads you will deploy after the POC. By the end of the POC configuration as described in this book, you will have deployed only a handful of VMs onto the compute cluster, so you probably won't exhaust your resources, and you should have headroom to add workloads as you see fit.

With this alternative minimal configuration, the requirements for the networking configuration drop to 14 10-Gbps network ports and cables and 16 1-Gbps network adapters. In the recommended hardware configuration, you need multiple switches for redundancy. However, it's acceptable if you only have one switch that supports both 1 Gbps and 10 Gbps. Alternatively, you can have one switch for each speed as long as they can be connected together. You will certainly lack redundancy, but the functionality should let you proceed with the POC configuration.

NOTE All of the steps in the book have been written based on the recommended hardware configuration. If you choose to use alternative hardware or align with the minimal hardware configuration, some steps may not apply.

Architecture

When you've understood and completed the hardware configuration, it's worthwhile to understand the overall architecture of this POC configuration before you take a deeper look at each of the core building blocks. The architecture is illustrated in Figure 1-1.

As you can see from Figure 1-1, several key building blocks make up the POC configuration. This chapter discusses each of these building blocks from a design perspective before subsequent chapters move on to the implementation of those building blocks.

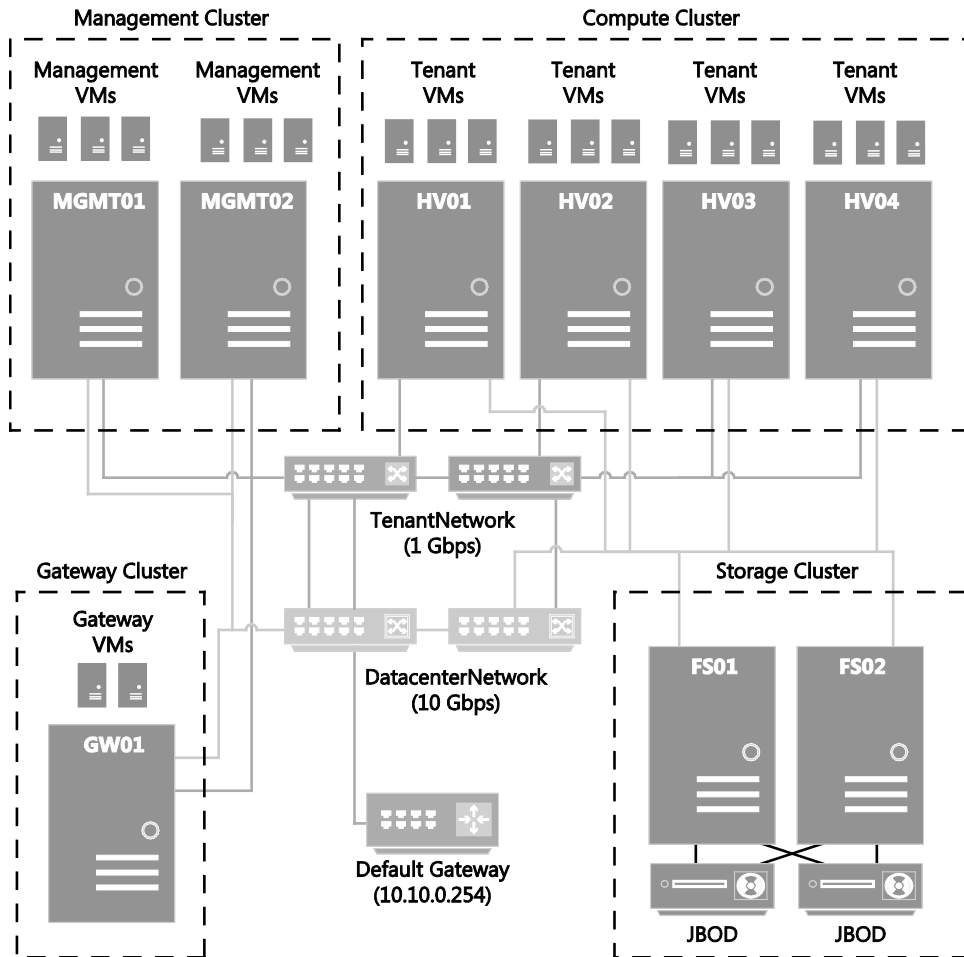


FIGURE 1-1 An architectural representation of the networks and clusters in the POC environment

At a high level, there are two core networks: the TenantNetwork, which runs at a speed of 1-Gbps, and the DatacenterNetwork, which runs at a speed of 10-Gbps. All of the physical servers are connected to the DatacenterNetwork. The storage nodes, labelled FS01 and FS02 in the figure, do not require a connection to the TenantNetwork, apart from a connection for their respective BMCs.

Two servers are dedicated as management nodes and are labelled MGMT01 and MGMT02. These will be the first Hyper-V hosts that you deploy. On these hosts, you'll create several VMs that will form the basis of your management backbone for the rest of the POC configuration. These management VMs will provide valuable directory services, operating system deployment and patching, and VM management functionality. The VMs will be spread across the two management nodes for effective load balancing and a greater level of redundancy.

When these are configured, you'll use the management infrastructure to deploy a new storage cluster, more specifically a Windows Server SOFS cluster that will provide the high-performance, resilient, and robust shared storage required to support the key workloads within the POC configuration. Because no VMs will technically run on either FS01 or FS02, these servers do not need to have a connection to the TenantNetwork, so no 1-Gbps network adapters are required.

Both of the file servers, FS01 and FS02, will be connected to a pair of JBOD enclosures by means of industry-standard 6-Gbps SAS cables. The JBOD enclosures will present their raw storage through to the Windows Server file servers. Through the power of software, Windows Server will transform them into the high-performance, resilient, robust SOFS. All subsequent nodes will use this shared storage for the deployment of their VM workloads from that point forward. The VMs that already existed on the management cluster, now residing on local storage on MGMT01 and MGMT02, will have their underlying virtual hard disks and configuration files live migrated onto the shared storage without incurring any downtime to the workloads. This provides a greater level of redundancy and availability for those workloads.

With the storage deployed, you'll see four Hyper-V compute nodes, HV01 through HV04, configured as a single Hyper-V cluster. These nodes will also be centrally deployed and then configured by the management infrastructure. These nodes will harness the 1-Gbps TenantNetwork to allow VMs that reside on the compute nodes to access the network in addition to the 10-Gbps DatacenterNetwork for node management, storage access, and live migration. The VMs will access the TenantNetwork through a logical switch and will be isolated from other VMs where appropriate, using network virtualization, all of which you'll learn more about later.

Finally, you'll deploy a Windows Server Gateway, GW01, consisting of a single Hyper-V host running VMs that collectively provide the desired Windows Server Gateway functionality and will unlock the full capabilities of network virtualization for the POC configuration. Without the Windows Server Gateway, the VMs that reside in VM networks would not be able to communicate out of their VM networks to the Internet, for example.

Now that you understand the architecture at a high level, it's worthwhile to delve deeper into the core building blocks of management; logical networking; and storage, compute, and network virtualization to ensure you are fully up to speed with the subtleties of the configuration and the reasoning behind why we chose this particular direction.

Management

As mentioned earlier, the management infrastructure will ideally consist of two modern, dual-socket x86 servers, each with two 10-Gbps and four 1-Gbps network adapters and a BMC. These management nodes should also have an additional 250 GB of local storage for storing VMs. You'll notice there is no specific guidance for memory. Based on the configuration of the VMs that you'll be placing on these management nodes, 48 GB of physical memory is the

recommended minimum. The reason for 48 GB is that your management VMs will likely consume up to approximately 34 GB of physical memory based on the configuration guidance in the next chapter. During maintenance operations, all VMs (except domain controllers, which will be discussed later) will likely need to run on a single physical management node. It's safe to plan for hosting all VMs, and thus approximately 34 GB of memory, on a single management node. You could use lower memory values for the management VMs or harness the Hyper-V Dynamic Memory functionality to allow Hyper-V to efficiently manage the memory allocated to VMs between an administrator-defined minimum and maximum. Both of these options would reduce the amount of memory required for the management nodes if your hardware configuration is memory constrained. The reason for choosing 48 GB instead of a figure closer to 34 GB is that it's more likely that the combination of memory DIMMs allows you to reach the 48 GB figure, and it provides you with plenty of headroom to add more management VMs in the future.

That covers memory, but why the requirement for local storage? Well, it's a little bit of a chicken-and-egg situation. Ideally, when you deploy VMs, you'll deploy them in such a way that their virtual hard disks and relevant configuration files reside on redundant, resilient, shared storage. However, at the start of this POC configuration, the shared storage doesn't exist yet. Some infrastructure—specifically domain controllers, as a minimum—need to be deployed first, so the local storage on the management nodes allows you to safely deploy the domain controller VMs, DC01 and DC02, without any requirement for shared storage. With the domain controllers deployed, you could proceed to manually install Windows Server on the physical file server nodes, FS01 and FS02, and configure a Windows Server SOFS cluster. However, it's a significant number of manual steps and doesn't expose you to some of the core System Center capabilities for fabric management that we believe are very important to learn. That's why you'll be deploying all of the management VMs onto the local storage of the management nodes. Once these are configured, you'll use the management infrastructure to automate the deployment of the shared storage. The management VMs' respective virtual hard disks and configuration files will then be migrated to the shared storage.

The final management storage question is whether you need 250 GB. It's unlikely. Each VM will be configured with a dynamic (thin-provisioned) virtual hard disk, most likely consuming 15 GB to 20 GB of physical capacity and thus spreading the management VMs evenly across MGMT01 and MGMT02, it's likely you'll need less than 150 GB per node. You will, however, also use some of that spare capacity to store important installation files, software, ISO files, and more, but even those additions shouldn't consume much more than another 15 GB or so.

With the hardware configuration finalized, you'll quickly move on to deployment. You'll notice in the early stages of the next chapter that the deployment of the management nodes is a manual process. Unfortunately, there was no easy way around this, at least from a documentation perspective. For a POC, we couldn't assume that all readers would have an already-configured environment with an operating system deployment mechanism and an existing domain for authentication and so on, that this POC configuration would live in. Thus, we decided to start from a blank canvas, with a completely fresh domain and forest and no

pre-existing capabilities to accelerate certain manual tasks. This means that at least a few of the core building blocks require some manual deployment. Fortunately, we've front-loaded the majority of the manual configuration, specifically the configuration of the management nodes, into Chapter 2. "Deploying the management cluster." This allows a much greater level of automation in the subsequent chapters as you allow the management infrastructure to streamline the ongoing deployment and configuration.

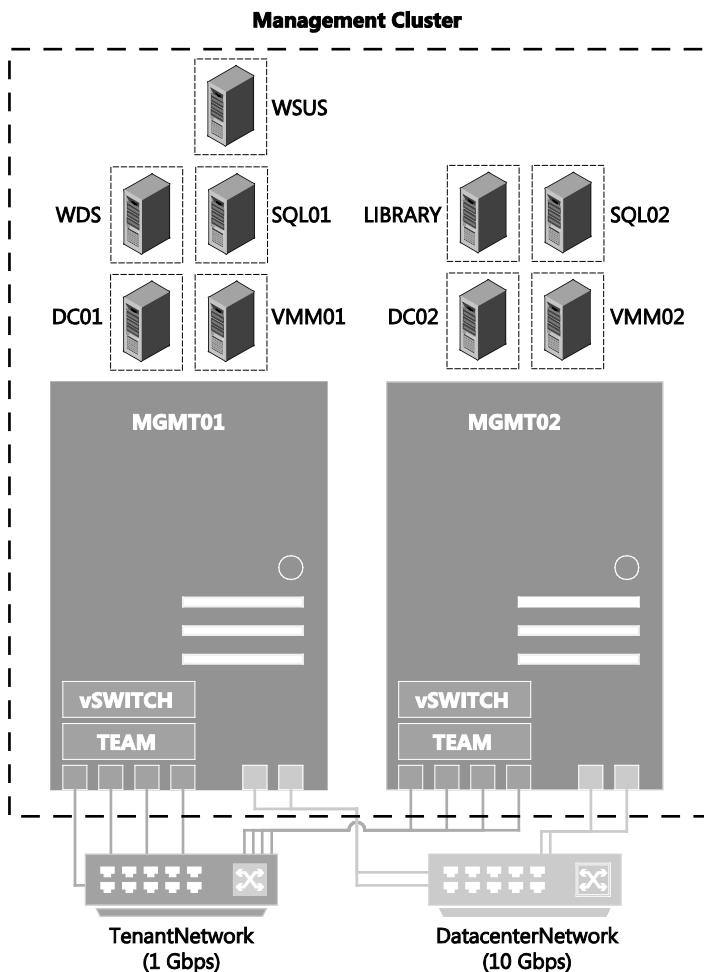


FIGURE 1-2 An architectural representation of the management infrastructure

With your management nodes deployed by means of your chosen operating system deployment mechanism, you'll move on to enabling the Hyper-V role, transforming the standalone physical Windows Server instances into a pair of powerful virtualization hosts. Before you create your management VMs, you'll need to manually configure some of the core Hyper-V settings. The most important of these settings, at least at this stage, is the creation of

a virtual switch, which will allow your management VMs to communicate over the physical TenantNetwork.

The Hyper-V virtual switch, or vSwitch, is a Layer 2 (L2) virtual network switch that provides programmatically managed and extensible capabilities to connect VMs to the physical network. The vSwitch provides policy enforcement for security, isolation, and service levels. With support for Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers, the Hyper-V vSwitch allows for third-party extensible plug-ins that can provide enhanced networking and security capabilities. In this POC configuration, you won't be exploring any of these switch extensions. However, going forward, if you wish to evaluate technologies such as the Cisco Nexus 1000V, Snine's Security Manager, NEC's PF1000 ProgrammableFlow virtual switch, or InMon Corp.'s sFlow switch extension, it's important to know that all of these technologies integrate through this extensible Hyper-V vSwitch.

With the Hyper-V vSwitch defined, you'll create management VMs that will provide the management backbone for the rest of the infrastructure. In Chapter 2, you'll deploy a total of eight management VMs, as follows:

- **Domain controllers (DC01 and DC02)** These VMs will be configured as a primary and secondary set of domain controllers, providing Active Directory Domain Services for a new domain, contoso.com, and forest. Both the domain and forest will be configured to operate at a Windows Server 2012 R2 functional level. The domain controllers will be distributed across MGMT01 and MGMT02. In addition to the Active Directory Domain Services, the domain controllers will both serve as Domain Name System (DNS) servers, providing name resolution for the contoso.com domain.
- **WDS** The WDS VM will be configured as a single WDS server and will enable the subsequent deployment of Windows operating systems over the network. This removes the requirement for DVD or USB deployment. Later, the WDS server will be integrated with System Center Virtual Machine Manager to orchestrate the deployment of the storage and compute nodes.
- **WSUS** The WSUS VM is not necessarily essential for a POC configuration, but it will allow you to centrally manage and distribute updates to the other machines within the contoso.com POC configuration. This WSUS server will download updates for operating systems and applications to a central source and deploy them to target servers across the local network, ensuring each server does not have to download updates from the Internet directly. Having WSUS configured will help to ensure that servers and applications are fully up to date with both security and feature-related patches and fixes.
- **SQL Server database services (SQL01 and SQL02)** The SQL Server VMs will provide the necessary database services for the System Center Virtual Machine Manager deployment. Without SQL Server, System Center Virtual Machine Manager cannot be installed and used, so SQL Server plays a pivotally important role within the management backbone. With that in mind, the two SQL Server VMs will be distributed across the two management nodes and configured as a SQL Server

AlwaysOn Availability Group. The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. The use of AlwaysOn Availability Groups in this POC configuration will maximize the availability of the System Center Virtual Machine Manager database, without the need, in this case, for shared storage.

- **System Center Virtual Machine Manager servers (VMM01 and VMM02)** The System Center Virtual Machine Manager VMs will host the Virtual Machine Manager component of the System Center product. System Center Virtual Machine Manager is a powerful and comprehensive management solution for the virtualized datacenter, enabling you to configure and manage your virtualization host, networking, and storage resources to create and deploy VMs and services to private clouds. In addition, System Center Virtual Machine Manager unlocks the power of network virtualization for the Hyper-V platform. You'll use System Center Virtual Machine Manager to perform the vast majority of the management of the infrastructure. Because it plays such an important role in the management backbone, it's imperative that it is as resilient and redundant as possible. With that in mind, System Center Virtual Machine Manager supports deployment in a highly available configuration, so VMM01 and VMM02 will be distributed across the underlying management nodes and across the two VMs that form a Windows Server failover cluster created to host the highly available System Center Virtual Machine Manager management service. This ensures that if one of the System Center Virtual Machine Manager management VMs suffers an outage, the other will take over to assure that you always have management control. You'll explore many of the features of System Center Virtual Machine Manager as part of the POC configuration.
- **Library services (LIBRARY)** Although you will not deploy library services until later in the POC configuration process, this library VM plays an important, integrated role with System Center Virtual Machine Manager. It enables the centralized storage of key System Center Virtual Machine Manager artifacts that System Center Virtual Machine Manager will use for deployment of Hyper-V hosts and VMs.

With the management VMs deployed, you'll use the management backbone that is now established to deploy the remaining network, storage, and compute infrastructure. With the shared storage deployed, you'll be able to transform your standalone management nodes, MGMT01 and MGMT02, into a more resilient, robust management cluster by using the powerful failover cluster capabilities that are built into Windows Server 2012 R2. By transforming the standalone nodes into a failover cluster and migrating the VM workloads—specifically their virtual hard disks and configuration files—to the shared storage, you will increase the levels of redundancy and resiliency considerably. This is of incredible importance for the management infrastructure.

The key benefit of using a failover cluster, rather than leaving the management nodes as standalone servers, is automated failover and increased redundancy. With the management VMs running on a failover cluster and their virtual hard disks and configuration files on shared

storage, if a management node fails, the management VMs on that failed node will automatically restart on the other management node within the failover cluster. This will take place without any intervention from you. This extra level of resilience and redundancy is important in ensuring the highest levels of availability for the infrastructure.

As mentioned in the discussion of the required memory for each management node, during a maintenance window, all VMs will run on a single management node while the other is patched and potentially rebooted. The exception is the domain controllers. When your shared storage is configured and deployed and all of the configuration files and virtual hard disks for the WSUS, WDS, SQL Server, and System Center Virtual Machine Manager servers have been migrated onto the shared storage, you'll leave the domain controllers, DC01 and DC02, running on the local storage of MGMT01 and MGMT02, respectively.

Madness, you may think. You want the domain controllers to be as resilient and redundant as possible, so storing their virtual hard disks and configuration files on shared storage makes perfect sense. There is, however, method to the madness. If you were to store the domain controllers on the shared storage and an outage occurred with the shared storage, the domain services would be unavailable. The unavailability of the domain services would make bringing the rest of the infrastructure back online potentially more difficult due to a variety of circular dependencies. Leaving the virtual hard disks and configuration files for DC01 and DC02, your domain controllers, on the local storage of MGMT01 and MGMT02, respectively, reduces the number of potential complications and simplifies troubleshooting in the event of an outage. You can configure DC01 and DC02 to automatically start when MGMT01 and MGMT02 start up. Because they reside on local storage, they have no reliance on the shared storage or compute infrastructures to be online and ready.

Logical networking

With the management cluster configured and the management VMs providing a robust control plane, you can begin to use System Center Virtual Machine Manager to design and architect the logical networking configuration that you will be using throughout the rest of the deployment.

This POC configuration has two core networks: one that runs at a speed of 1 Gbps and another that runs at a speed of 10 Gbps. The 1-Gbps network will be used predominantly for the VMs in the environment and the BMCs for each physical server. Note that some models of BMC operate at 1-Gbps speed and others may only operate at 100 Mbps. Either way, we will refer to the network as having 1-Gbps speed. The other core network, the 10-Gbps network, will provide multiple functions for the POC configuration. It will carry storage traffic between the storage nodes (FS01 and FS02) and the management, compute, and gateway nodes. In addition, it will be used for live migration traffic as VMs are migrated between the respective management and compute nodes. In addition, this 10-Gbps network will carry management traffic. These two core networks will each be mapped into a System Center Virtual Machine Manager construct known as a logical network.

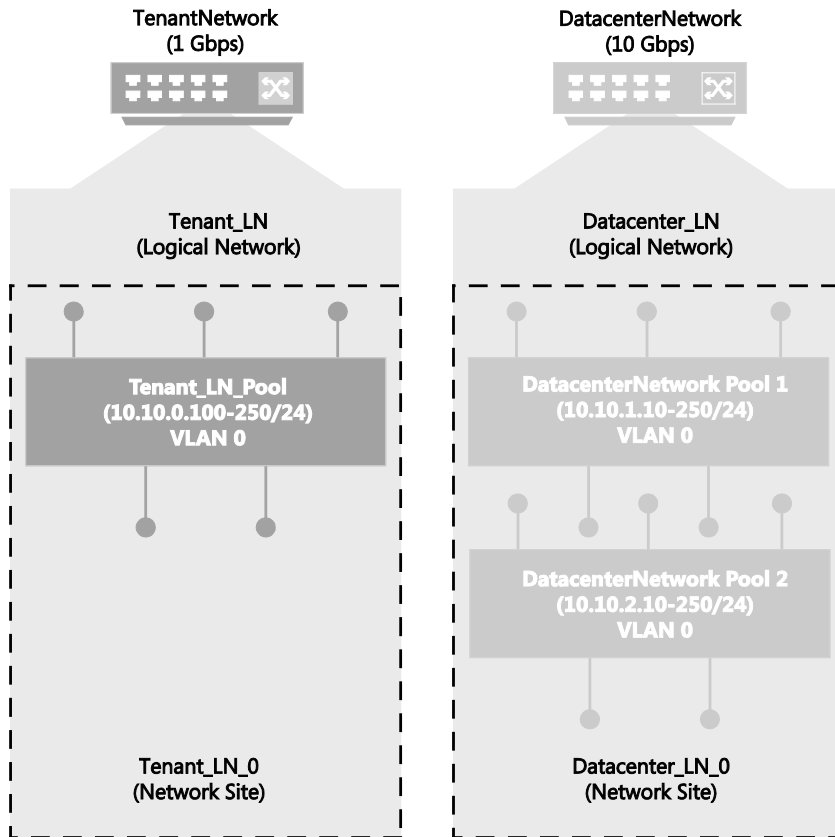


FIGURE 1-3 An architectural representation of the Tenant_LN and Datacenter_LN logical networks

At a high level, logical networks are named networks that serve particular functions in your environment. For example, the Backend, Frontend, or Backup network. In this POC configuration, you will ultimately create two logical networks. As shown in Figure 1-3, the 1-Gbps TenantNetwork is labelled as Tenant_LN and the 10-Gbps DatacenterNetwork is labelled as Datacenter_LN. The functions of these two logical networks is simple: The Tenant_LN logical network will ultimately carry tenant traffic (from tenant VMs) but also will be used for the management VMs and the BMCs. The Datacenter_LN logical network will be used for infrastructure traffic, as mentioned earlier, which includes storage, live migration, and management traffic.

A logical network is a container for network sites (also called logical network definitions) and for IP subnet information, virtual local area network (VLAN) information, or both. Host groups in System Center Virtual Machine Manager can be associated with a network site, and IP address pools can be assigned to subnets within the logical network. In this small POC configuration, you will configure an individual network site for both the Tenant_LN logical network and the Datacenter_LN logical network (Tenant_LN_0 and Datacenter_LN_0, respectively). Within these sites, you will configure the appropriate IP subnets and IP pools that

System Center Virtual Machine Manager will use to distribute IP addresses to physical machines or VMs that reside on that network.

Notice in Figure 1-3 that the Tenant_LN logical network has a single subnet, 10.10.0.0/24, and corresponding IP pool from 10.10.0.100 to 10.10.0.250. System Center Virtual Machine Manager will distribute IP addresses from that range to new VMs later in the POC configuration process. Also notice that the Datacenter_LN logical network has two subnets, 10.10.1.0/24 and 10.10.2.0/24, each with a corresponding IP pool. The reason for creating two subnets for the Datacenter_LN logical network, each with a corresponding IP pool, is to embrace a key capability known as Server Message Block (SMB) multichannel, which you'll learn more about later. Essentially, when System Center Virtual Machine Manager configures your storage, compute, and gateway nodes, one of the 10-Gbps network adapters on each node is given an IP address on the 10.10.1.0/24 subnet and the other 10-Gbps network adapter on each node is given an IP address on the 10.10.2.0/24 subnet. Separating these ports on different subnets allows each of these nodes to take advantage of SMB multichannel.

It's important to note that although the Tenant_LN and Datacenter_LN logical networks have different subnets, they are routable. Therefore, your management VMs, which will ultimately reside on the Tenant_LN logical network and have an IP address in the range of 10.10.0.0/24, can still communicate and manage the nodes that reside on the 10.10.1.0/24 and 10.10.2.0/24 subnets. This routing would typically be configured on your switches. In a production environment, it's also important to note that these logical networks could be isolated from one another by means of a VLAN. Again, you would configure the appropriate routing at the switch level to route between the logical networks where appropriate, for instance, where the management VMs on the Tenant_LN logical network need to communicate with physical hosts residing on the Datacenter_LN logical network. The aim of this POC configuration is to keep the logical networking as simple and as flat as possible, which is why VLAN 0 has been used throughout.

By configuring all of your logical networks, network sites, IP subnets, and IP pools before deploying your storage, compute, and gateway nodes, you are putting steps in place to significantly streamline the deployment and configuration process of those respective nodes later. As you'll learn shortly, when you use System Center Virtual Machine Manager to deploy your storage, compute, and gateway nodes, it will apply the relevant logical network settings in a predictable and controlled manner at deployment time. This ensures that your storage, compute, and gateway nodes integrate seamlessly with your existing setup as quickly as possible, without you having to perform any network-related configuration post-deployment.

Logical networks enable you to model your network design and requirements effectively. You will take advantage of several other networking constructs within System Center Virtual Machine Manager as part of this POC configuration. One of those key constructs is the logical switch. You'll learn considerably more about the logical switch and its core building blocks in Chapter 3, "Configuring network infrastructure". At a high level, logical switches allow VMs to communicate out through the physical adapters and network adapter teams configured on a Hyper-V host. Because System Center Virtual Machine Manager is not available when you first

deploy your management nodes, you do not have the ability to use the logical switch. It is a System Center Virtual Machine Manager-only construct, which is why you will deploy a standard Hyper-V vSwitch to those management nodes. A logical switch still harnesses those same Hyper-V vSwitch capabilities but offers greater management granularity and control from within System Center Virtual Machine Manager.

Logical switches can also enforce standard configurations on Hyper-V host adapters and NIC teams to prevent workloads from experiencing downtime due to misconfigurations. For instance, if someone logged on to the Hyper-V host directly and made a configuration change to the vSwitch, System Center Virtual Machine Manager could detect and remediate this to bring the vSwitch back in line with its original deployment configuration. This is a capability unique to a logical switch.

As part of the POC, you will create a single logical switch and all of the core building blocks that it consists of. Upon deployment of your compute and gateway nodes, System Center Virtual Machine Manager will automatically aggregate the 1-Gbps network adapters in each respective host into a network team and bind the logical switch to that team. When your nodes have completed the deployment process, they will all have an identical, standardized instance of the logical switch across each node, ensuring efficient ongoing management and consistent operations.

Storage

If we'd looked at a traditional datacenter when designing the storage configuration for this POC, we'd most likely have landed on the use of either an iSCSI or Fibre Channel-based storage area network or SAN. iSCSI and Fibre Channel SANs have been the storage deployment of choice within the datacenter for a considerable time, and for some datacenters, that will continue for years to come.

Have you ever taken the time to think just what a SAN is, what's inside the chassis, and what it provides to the infrastructure?

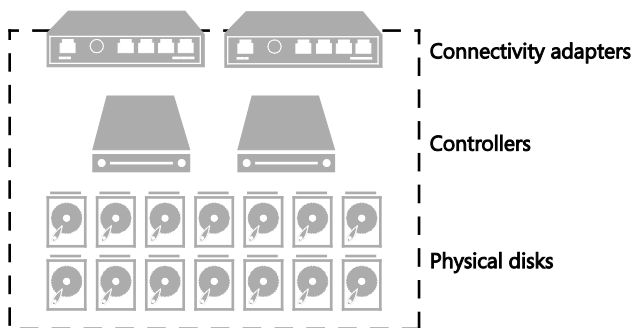


FIGURE 1-4 Conceptual layout of SAN storage

If you were to peek inside the chassis, as shown in Figure 1-4, you'd very likely find three key elements. First, you'd have physical disks. These could be exclusively Serial Attached SCSI (SAS) or Serial Advanced Technology Attachment- (SATA-) based HDDs with large capacities and speeds that could vary from 7,200 RPM up to 15,000 RPM. In addition, within more modern SANs, you may find SSDs with smaller capacities but significantly higher performance. The SSDs are typically in addition to the HDDs. In some SANs, you'll find only SSDs, which provides incredible performance but at significantly higher cost.

To transform that raw capacity into usable storage, you'll find the controller, or in most cases, controllers. Two or more controllers provide the brains of the SAN, transforming the physical disks into aggregated pools of redundant storage in the form of aggregates, volumes, or logical unit numbers (LUNs), depending on your SAN specifics. If you were to look inside one of these controllers, you'd likely find a modern x86 processor, memory, and perhaps some sort of local storage, on which perhaps a bespoke, heavily customized operating system is running. With two or more controllers, you have redundancy against failure, but you also have redundancy at times of maintenance and firmware updates. Aside from aggregating the storage into usable striped or mirrored volumes, these controllers provide advanced features such as deduplication, thin provisioning, rapid cloning, snapshots, storage tiering, and more. It's these advanced features, combined with the reliability and redundancy associated with SANs, that drive so many organizations to choose them as their preferred infrastructure storage.

The final element within the SAN is the connectivity. How do you present the usable storage that you, the administrator, have created from the raw physical disks? With many SANs, you have a choice. iSCSI connectivity may be the preference since it uses the Ethernet network as transport and can support both 1-Gbps and 10-Gbps connectivity. Alternatively, perhaps Fibre Channel is the preferred connection option because it has always been perceived as a higher performing option. However, Fibre Channel brings a requirement for specific Fibre Channel switches and host bus adapters (HBAs), and this typically drives the cost of a Fibre Channel deployment higher than an equivalent iSCSI-based deployment. It's typical for storage vendors to offer an either/or approach to connectivity, but some SANs do enable a network and Fibre Channel connectivity offering in the same chassis.

It's interesting to understand the core building blocks of a SAN. However, for this POC configuration, you'll be taking a different approach. The design of the storage infrastructure for this POC shares many similarities with the SAN-breakdown you just read, but this Microsoft software-defined storage (SDS) infrastructure is constructed from low-cost, high-volume hardware, transformed with the power of software.

At first glance, the conceptual layout, as shown in Figure 1-5, of the Microsoft SDS infrastructure looks similar to that of the SAN. While there are similarities, there are also a number of differences.




FIGURE 1-5 Key storage elements in the server and JBOD infrastructure

First, you have the physical disks. Again, this is similar to the SAN. These could be SAS- or SATA-based HDDs or SSDs. This time, instead of being contained within a SAN, they are contained with an external JBOD enclosure. This configuration will be using two 12-bay JBOD enclosures, for a total of 24 disk drives. When you deploy in a production environment, the best practice is to have SSDs as at least 20 percent of the drives and HDDs as the remaining 80 percent. An important best practice is that they be SAS-based HDDs and SSDs because these offer greater performance and reliability than SATA equivalents.

Within each JBOD enclosure are typically multiple SAS interfaces. These will be used to connect the respective JBOD enclosures to the controllers, which, in the Microsoft SDS solution and in this POC configuration, will be a pair of Windows Server 2012 R2 servers (FS01 and FS02). The SAS interfaces should each be multimaster- or multiserver-capable. The drives within the enclosures must support port association. Multiport drives must provide symmetric access, and drives must provide persistent reservations.

To identify disks by slot and take advantage of the enclosure's failure and identify/locate lights, the JBOD enclosures must support SCSI Enclosure Services (SES) version 3. In addition, the enclosure must provide direct access to the drives that are housed within the enclosure. It's important not to apply any abstraction to the drives such as grouping drives into a RAID array. The drives should be presented raw to the Windows Server 2012 R2 controller nodes. Finally, the JBOD enclosure must provide multiple SAS interfaces to the drives and to the controllers to provide the highest levels of reliability and redundancy.

For a specific list of JBOD enclosures that are supported with Windows Server 2012 R2, refer to the Windows Server Catalog at

<http://windowsservercatalog.com/results.aspx?&chtext=&cstext=&csstext=&chbtext=&bCatID=1645&cpID=0&avc=10&ava=80&avq=0&OR=1&PGS=25&ready=0>.

As mentioned earlier, the JBOD enclosures will be connected to the Windows Server 2012 R2 controller nodes by means of industry-standard SAS cables, at either 6-Gbps or 12-Gbps

speed, depending on your specific JBOD enclosure and the SAS interfaces within your controller nodes.

The SAS interfaces within your Windows Server 2012 R2 controller nodes should be at least 6 Gbps. If you are using RAID adapters instead of SAS HBAs, they must be in non-RAID mode with all RAID functionality disabled. Such adapters must not abstract the physical disks or cache data or obscure any attached devices, including enclosure services provided by the JBOD enclosures.

Continuing through the layers of the conceptual architecture, you have the controllers. In the SAN configuration that was discussed earlier, multiple controllers were available for redundancy. In addition, these controllers provided the intelligence to wrapper the physical disks with advanced features and functionality. Using multiple Windows Server 2012 R2 servers as the controller nodes provides similar capabilities as part of the Microsoft SDS solution.

With the raw physical drives presented from the JBOD enclosures through to both FS01 and FS02, you'll use the advanced functionality of Windows Server 2012 R2 to greatly simplify storage management and storage virtualization while creating reliable, highly available storage.

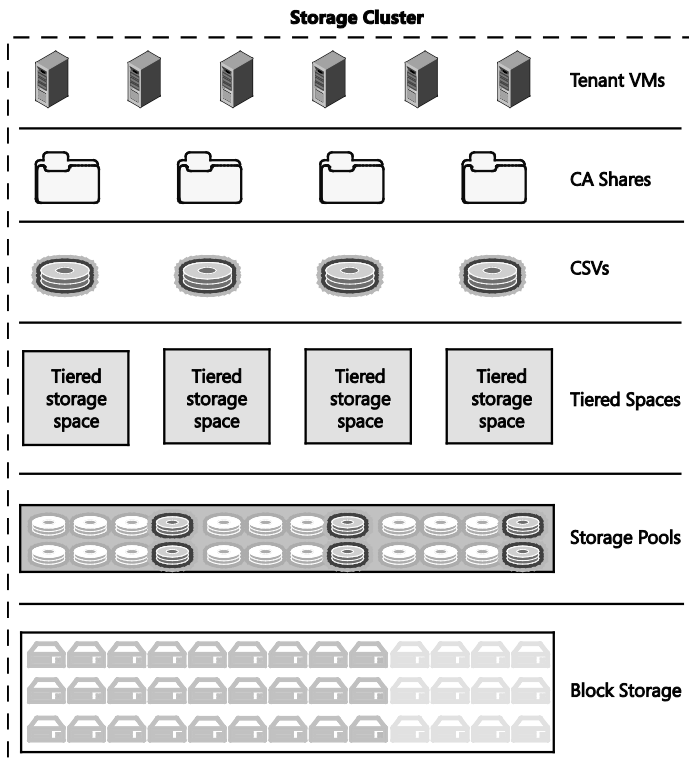


FIGURE 1-6 Microsoft Software Defined Storage layers

You'll transform these raw physical drives into a single storage pool. As shown in Figure 1-6, these storage pools are virtualized units of administration that are aggregates of underlying physical disk drives. They aggregate the capacity but also allow for elastic capacity expansion if you add more drives, and delegated administration. When combined with Failover Clustering, the clustered pool can be used by any node in the file server cluster and fail over within the failover cluster. This is useful to automatically recover from failures, as well as for load balancing workloads.

From there, you'll slice that pool into a number of storage spaces. Storage spaces are virtual disks with attributes that include a desired level of resiliency, fixed or thin provisioning, automatic or controlled allocation of heterogeneous classes of storage, and precise administrative control. In the traditional SAN world, this would equate to creating a LUN or volume from your aggregated set of disks.

Windows Server 2012 R2 supports aggregating HDDs and SSDs into a single, logical disk, or storage space, allowing configurations that seamlessly provide both the high capacity of HDDs and the high performance and low latency of SSDs. HDDs and SSDs are defined as different tiers as part of the same storage space. Even though there are two distinct classes of hardware, a tiered space is presented as a seamless, unified disk. File systems will track how frequently data located on the tiered space is accessed and will periodically automatically move frequently accessed data to the portion of the tiered space backed by SSDs. Infrequently accessed data will similarly be moved to the portion of the tiered space backed by HDDs. Frequently accessed data will be stored on SSDs, providing low latency access to applications. Infrequently accessed data will be stored on low-cost, high-capacity HDDs, ensuring that data placement is responsive to varying workload conditions. In Windows Server 2012 R2, storage tiering is compatible only with simple or mirrored storage spaces.

A single storage space and the file system that resides on top of the storage space are controlled or resident on exactly one storage node at a time. Using Failover Clustering, the storage space can move to a different storage node within the cluster. The movement of a storage space will occur if the original storage node is no longer functional. The storage space can also be moved if necessary to achieve a balanced load within the storage nodes of the cluster. Layering Cluster Shared Volumes (CSVs) on top of the storage space/file system combination creates a clustered file system that all services on any storage node in the cluster can then access. If a storage space is moved anywhere within the cluster, CSV will transparently redirect all I/O from the service that is using the file system or storage space to the new storage node that the file system or storage space resides on.

Ultimately, by transforming your Windows Server 2012 R2 controller nodes FS01 and FS02 into a file server cluster and with that configuring clustered storage pools, spaces, and CSVs, you're in a position to deploy the SOFS feature. SOFS is a feature that is designed to provide scale-out file shares that are continuously available for file-based server application storage. SOFS shares provide the ability to share the same folder from multiple nodes of the same cluster and are therefore perfect to store the Hyper-V VMs that will run on your compute, management, and gateway nodes.

This configuration will use two file server nodes as the controllers for the SDS infrastructure. However, up to eight nodes are supported, which unlocks significant scalability and provides the desired redundancy in case of a node failure. You'll also find many other benefits to using the SOFS:

- **Active-Active file shares** All cluster nodes can accept and serve SMB client requests. Because the file share content is accessible through all cluster nodes simultaneously, SMB 3.0 clusters and clients cooperate to provide transparent failover to alternative cluster nodes during planned maintenance and unplanned failures with service interruption.
- **Increased bandwidth** The maximum share bandwidth is the total bandwidth of all file server cluster nodes. Unlike in previous versions of Windows Server, the total bandwidth is no longer limited to the bandwidth of a single cluster node. Instead, the capability of the backing storage system defines the constraints. You can increase the total bandwidth by adding nodes.
- **CHKDSK with zero downtime** CHKDSK in Windows Server 2012 R2 is significantly enhanced to dramatically shorten the time a file system is offline for repair. CSVs take this one step further by eliminating the offline phase. A CSV File System (CSVFS) can use CHKDSK without impacting applications with open handles on the file system.
- **Clustered Shared Volume cache** CSVs in Windows Server 2012 R2 provide support for a read cache, which can significantly improve performance in certain scenarios, such as in Virtual Desktop Infrastructure (VDI).
- **Automatic rebalancing of SOFS clients** In Windows Server 2012 R2, automatic rebalancing improves scalability and manageability for scale-out file servers. SMB client connections are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share. This improves efficiency by reducing redirection traffic between file server nodes. Clients are redirected following an initial connection and when cluster storage is reconfigured.

The final layer in the conceptual model is connectivity. This POC configuration will harness the 10-Gbps DatacenterNetwork as the transport to connect the relevant management, compute, and gateway nodes to the SOFS. Harnessing Ethernet-based connectivity eases integration and streamlines management.

Harnessing the two 10-Gbps network adapters means each node will have up to 20 Gbps of bandwidth. However, this isn't dedicated to just storage. You'll learn in later chapters how the use of quality of service (QoS) controls can help to allocate percentages of this total bandwidth for different traffic types, including management, live migration, and storage traffic.

For storage connectivity from the compute, management, and gateway nodes, you will harness the power of SMB 3.0 to deliver a high-performance, redundant, and simple-to-manage storage network infrastructure. By using SMB 3.0 as the transport for storage traffic, you'll benefit from several key capabilities:

- **SMB Transparent Failover** SMB Transparent Failover enables administrators to configure file shares in a Windows Server failover cluster configuration to be continuously available. This means that when you perform hardware or software maintenance on nodes in a clustered file server, there is no interruption to the server applications storing data on these file shares. Also, if a hardware or software failure occurs on a cluster node, SMB clients transparently reconnect to another cluster node without interrupting server applications that are storing data on these file shares. When you use an SMB 3.0 scale-out file share, SMB 3.0 transparent failover also allows you to redirect a server application node to a different storage node to facilitate better load balancing.
- **SMB Multichannel** SMB Multichannel enables aggregation of network bandwidth and network fault tolerance if multiple paths are available between the SMB 3.0 client and the SMB 3.0 server. This allows server applications to take full advantage of all available network bandwidth and to be resilient in case of a network failure. This is one of the key reasons for configuring the DatacenterNetwork with two subnets, which provides two paths so that SMB Multichannel can function correctly.
- **SMB Direct** The SMB 3.0 protocol in Windows Server 2012 R2 includes support for RDMA network adapters, which allows storage-performance capabilities that surpass those of Fibre Channel. RDMA-capable network adapters make this performance possible by operating at full speed with very low latency because of their ability to perform zero-copy data transfers for both write and read operations while preserving full security. This capability requires advanced functionality in the network adapter hardware. Using this capability, SMB 3.0 can perform data transfers directly between the tenant application's memory and the storage node's memory, creating a zero-copy path through the adapter, across the network, and all the way to the storage device. In addition, because of the low overhead of the transfer, the CPU utilization is similar to or better than Fibre Channel HBAs, even when a conventional Ethernet network is used. This capability is especially useful for read-intensive and write-intensive workloads, such as Hyper-V or Microsoft SQL Server, and results in remote file-server performance that is comparable to local storage. You will learn more about RDMA in Chapter 5, "Configuring compute infrastructure," as you will perform specific configurations when your storage and compute infrastructure is established.

Now that you understand the architecture of the software-defined storage solution, specifically, the SOFS, it's important to understand how it will be deployed. You could manually deploy it in the same manner that you will deploy the management cluster. However, with the management VMs up and running, it makes sense to streamline and considerably accelerate the deployment and configuration process by using System Center Virtual Machine Manager to deploy and configure the storage infrastructure.

System Center Virtual Machine Manager automates many of the key deployment and configuration tasks, leaving only minimal additional configuration, as you'll find in Chapter 4, "Configuring storage infrastructure." You'll use the bare-metal deployment capabilities of

System Center Virtual Machine Manager to deploy a standardized image across to the FS01 and FS02 servers. From there, System Center Virtual Machine Manager will handle the rest, from configuring the Windows Server operating system, joining the domain, applying the relevant network configuration, and even creating the SOFS cluster. Then, post-deployment, System Center Virtual Machine Manager will ease you through creation of storage pools and the continuously available file shares, ensuring that your resilient, high-performance software-defined storage infrastructure is up and running quickly and efficiently.

Compute

With the logical networks and logical switch created and the resilient SDS deployed and operational, you'll use your management infrastructure to deploy additional Hyper-V nodes. These Hyper-V nodes will make up a compute cluster. This is the compute engine your tenant workloads will run on.

Your hardware for the compute nodes should consist of four modern, dual-socket x86 servers, each with two 10-Gbps and four 1-Gbps network adapters and a BMC. As with the management nodes before them, the 10-Gbps adapters attached to the Datacenter_LN logical network will connect the compute nodes to the SDS infrastructure. In addition, this logical network will also carry the management traffic and the traffic associated with live migration of VMs. The four 1-Gbps network adapters will connect to the Tenant_LN logical network and will be the first adapters to use the logical switch that you will define in Chapter 3.

One question that comes to mind is whether you really need four compute nodes. The answer is no; as described earlier, in the minimum configuration, two compute nodes may be sufficient for your needs. By the end of this book, you will have deployed only a handful of VMs onto your compute nodes. This small number of VMs could easily run on two, or even one compute node. However, you should have at least two compute nodes to assure that you can deploy a compute cluster. This increases resilience and redundancy while unlocking other interesting capabilities that you will explore in later chapters. Having four compute nodes will allow you to deploy additional workloads after the completion of the guidance in this book.

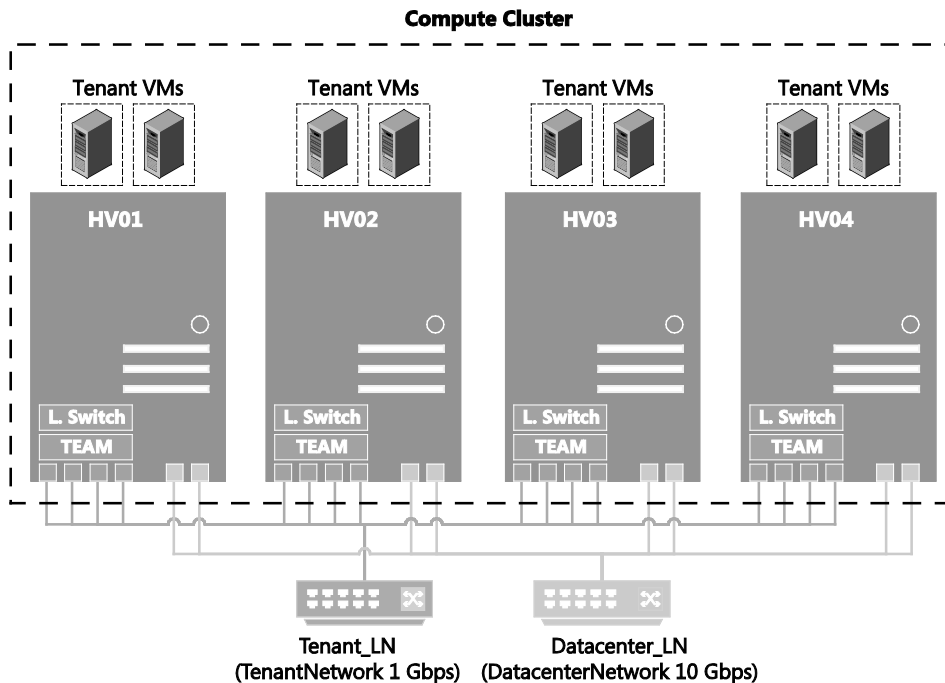


FIGURE 1-7 Logical architecture for the compute cluster

To deploy the compute nodes, you'll harness the same deployment capabilities that you'll use to deploy your SDS infrastructure—specifically, your SOFS. System Center Virtual Machine Manager will orchestrate the deployment process and deploy the appropriate image to the physical nodes. It will use the BMC for the initial communication and deployment steps. The node will have the Hyper-V role enabled automatically and have the networking configuration automatically applied based on your logical networking definitions you will create in Chapter 3. This means that each of the 10-Gbps network adapters will be assigned an IP address from the respective IP pools within the Datacenter_LN logical network, which in turn will unlock communication back to the management VMs. In addition to applying the Datacenter_LN settings, System Center Virtual Machine Manager will automatically transform the four 1-Gbps network adapters into a redundant network team and deploy the logical switch to that network team. This will simultaneously take place automatically across all four compute nodes. When the process completes, you will have four compute nodes, all with a standardized configuration. When you compare this to the manual approach you have to take to deploy your initial management nodes, you'll find the process incredibly streamlined and straightforward.

With the deployment completed, you need to perform a small number of extra steps before the compute cluster is operational. You'll walk through assigning the shared storage to the new compute cluster to ensure that you can deploy tenant workloads onto a resilient

infrastructure. You'll also enable features to optimize the compute cluster configuration and configure the live migration settings for the compute cluster.

As you may know, Hyper-V Live Migration moves running VMs from one physical server to another with no impact on VM availability to users. With Windows Server 2012 R2, you can take advantage of the power of faster networks to live migrate VMs faster than ever, and you can migrate multiple VMs simultaneously.

For larger-scale deployments, such as private cloud deployments or cloud hosting providers, Windows Server 2012 R2 introduced improvements in Live Migration that can reduce overhead on the network and CPU usage, in addition to reducing the amount of time for a live migration. You will use System Center Virtual Machine Manager to configure the appropriate Live Migration settings, but for reference, the three different types of live migration are:

- **TCP/IP** In a TCP/IP live migration, the memory of the VM is copied to the destination server over a TCP/IP connection. This is the same method that Windows Server 2012 Hyper-V uses.
- **Compression** A compression live migration uses spare Hyper-V node resources to compress the memory content of the VM that is being migrated. Then the memory content is copied to the destination server over a TCP/IP connection. This is the default Live Migration setting in Windows Server 2012 R2.
- **SMB** An SMB live migration copies the memory content of the VM to the destination server over an SMB 3.0 connection. Using Live Migration over SMB means that the live migration process can benefit from SMB capabilities such as SMB Multichannel and SMB Direct. This method is the fastest live migration of the three options.

To enable SMB Direct, you will need to configure the environment to support RDMA. You'll learn a significant amount about RDMA at the end of Chapter 5. If you have appropriate hardware, you'll put in place a number of configuration settings that let you considerably accelerate the performance across the Datacenter_LN logical network.

Network virtualization

By the end of Chapter 5, you'll have a fully operational infrastructure across management, storage, and compute clusters. You'll be ready to deploy your virtual workloads onto your compute cluster and to understand some of the additional features that Hyper-V and System Center Virtual Machine Manager can provide.

If you were to immediately begin deploying VMs onto your compute cluster, joining them to a virtual machine network that in turn was joined to the underlying Tenant_LN logical network, these new VMs would be assigned IP addresses from the Tenant_LN pool. They'd be in a position to communicate back to your management VMs and the other physical nodes in

the environment. This may be fine for a smaller environment. But if you want to isolate your management infrastructure from these new tenant workloads, you'll need to put controls in place to enable isolation.

Currently, VLANs are the mechanism that most organizations use to support isolation within a network environment. A VLAN uses explicit tagging (VLAN ID) in the Ethernet frame headers, and it relies on Ethernet switches to enforce isolation and restrict traffic to network nodes with the same VLAN ID. However, VLANs also have some disadvantages:

- Increased risk of an inadvertent outage due to cumbersome reconfiguration of production switches whenever VMs or isolation boundaries move in a dynamic datacenter.
- Limited scalability because 4,095 is the maximum number of VLANs, but typical switches support no more than 1,000 VLAN IDs.
- Restricted to residing within a single IP subnet, which limits the number of nodes within a single VLAN and limits the placement of VMs based on physical locations. Even though VLANs can be expanded across sites, the entire VLAN must be on the same subnet.

For a smaller environment, such as this POC configuration, the second bullet point above is unlikely ever to come into question. The maximum 4,095 VLANs is still a significant number, yet it's likely more advanced, and, thus, expensive switch infrastructure would be required to support that level of scale. But even support for 1,000 VLAN IDs is plenty for most organizations. The cumbersome reconfiguration, however, can impact small and large environments alike, so a simplification of this process, reducing the potential for inadvertent outages, would be welcome.

Just as detailed in the storage section earlier, when you're planning this POC configuration, you could take the approach to use networking technologies that are most common in the datacenter. However, this POC uses a SDS approach instead of a traditional SAN. Similarly, networking isolation will use network virtualization to handle isolation needs rather than using traditional VLANs.

Hyper-V Network Virtualization decouples virtual networks for tenant VMs from the physical network infrastructure. This provides benefits, including flexible workload placement, network isolation, and IP address re-use without VLANs. Hyper-V Network Virtualization decouples the tenants' virtual networks from the underlying physical network infrastructure, providing freedom for workload placements inside the datacenters. VM workload placement is no longer limited by the IP address assignment or VLAN isolation requirements of the physical network because it is enforced within Hyper-V hosts based on software-defined, multitenant virtualization policies.

You can now deploy VMs from different tenants with overlapping IP addresses on the same host server without requiring cumbersome VLAN configuration or violating the IP address hierarchy. This can streamline the migration of customer workloads into shared infrastructure as a service (IaaS) cloud service providers, allowing you to move those workloads without

modification. This includes leaving the VM IP addresses unchanged. For the cloud service provider, supporting numerous customers who want to extend their existing network address spaces to the shared IaaS datacenter is a complex exercise of configuring and maintaining isolated VLANs for each customer to ensure the coexistence of potentially overlapping address spaces. Hyper-V Network Virtualization makes it easier to support overlapping addresses and requires less network reconfiguration by the hosting provider.

In addition, cloud service providers can perform physical infrastructure maintenance and upgrades without causing downtime for customer workloads. With Hyper-V Network Virtualization, VMs on a specific host, rack, subnet, VLAN, or an entire cluster can be migrated without requiring a physical IP address change or major reconfiguration.

The use of network virtualization also enables live migration across subnets. Live migration of VM workloads traditionally has been limited to the same IP subnet or VLAN because crossing subnets required the VM's guest operating system to change its IP address. This address change breaks existing communication and disrupts the services running on the VM. With Hyper-V Network Virtualization, you can live migrate workloads from servers running in one subnet to servers running in a different subnet without changing the workload IP addresses. Hyper-V Network Virtualization ensures that VM location changes caused by live migration are updated and synchronized among hosts that have ongoing communication with the migrated VM.

Network virtualization also makes it easier to manage decoupled server and network administration. Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, and network administrators can focus on overall network infrastructure and traffic management. This lets datacenter server administrators deploy and migrate VMs without changing the IP addresses of the VMs. This reduces overhead because Hyper-V Network Virtualization allows VM placement to occur independently of network topology, which reduces the need for network administrators to be involved with placements that might change the isolation boundaries.

Network virtualization can significantly simplify the network and ultimately improve server and network resource utilization. The rigidity of VLANs and the dependency of VM placement on a physical network infrastructure results in overprovisioning and underutilization. Breaking the dependency increases flexibility of VM workload placement and can simplify the network management and improve server and network resource utilization. It's important to note that Hyper-V Network Virtualization still supports VLANs in the context of the physical datacenter. For example, a datacenter may want all Hyper-V Network Virtualization traffic to be on a specific VLAN.

Best of all, and one of the reasons you can deploy network virtualization in this POC configuration, is that it is compatible with existing infrastructure and emerging technology. It also provides for interoperability and ecosystem readiness. Hyper-V Network Virtualization supports multiple configurations for communication with existing resources such as cross-premise connectivity, SANs, non-virtualized resource access, and so on. Microsoft is committed

to working with ecosystem partners to support and enhance the experience of Hyper-V Network Virtualization performance, scalability, and manageability.

Now that you understand some of the benefits of network virtualization, how does it actually work? Well, in some ways, it helps to think of network virtualization in the context of server virtualization.

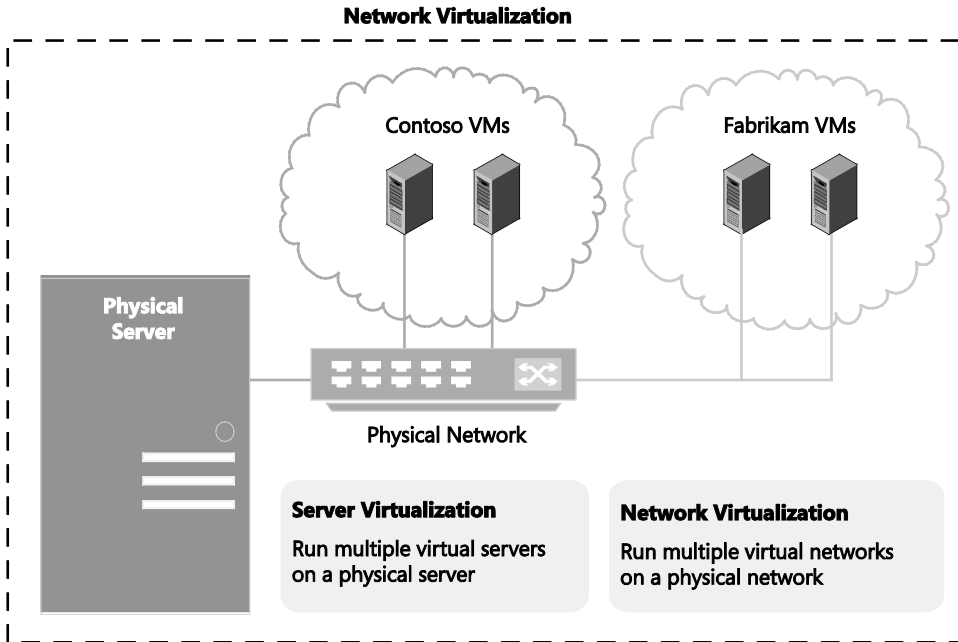


FIGURE 1-8 Logical architecture for network virtualization

As shown in Figure 1-8, server virtualization enables multiple server instances to run concurrently on a single physical host. Yet server instances are isolated from each other. Each VM essentially operates as if it is the only server running on the physical computer. Network virtualization provides a similar capability, in which multiple virtual network infrastructures run on the same physical network (potentially with overlapping IP addresses), and each virtual network infrastructure operates as if it is the only virtual network running on the shared network infrastructure.

With Hyper-V Network Virtualization, a customer or tenant is defined as the owner of a group of VMs that are deployed in a datacenter. A customer can be a corporation or enterprise in a multitenant public datacenter, or a division or business unit within a private datacenter. Each customer can have one or more VMs or VM networks in the datacenter, and each VM network consists of one or more virtual subnets.

- **VM network** Each VM network consists of one or more virtual subnets. A VM network forms an isolation boundary so that the VMs within a VM network can communicate with each other. As a result, virtual subnets in the same VM network

must not use overlapping IP address prefixes. Each VM network has a routing domain that identifies the VM network. The routing domain ID (RDID), which identifies the VM network, is assigned by datacenter administrators or datacenter management software such as System Center Virtual Machine Manager. The RDID is a Windows GUID, for example, {11111111-2222-3333-4444-000000000000}

- Virtual subnet** A virtual subnet implements the Layer 3 (L3) IP subnet semantics for the VMs in the same virtual subnet. The virtual subnet is a broadcast domain (similar to a VLAN). VMs in the same virtual subnet must use the same IP prefix. Each virtual subnet belongs to a single VM network (RDID), and it is assigned a unique Virtual Subnet ID (VSID). The VSID must be unique within the datacenter and is in the range 4096 to $2^{24}-2$.

A key advantage of the VM network and routing domain is that it allows customers to bring their network topologies to the cloud. Figure 1-9 shows an example where Contoso has two separate networks, the R&D Net and the Sales Net. Because these networks have different routing domain IDs, they cannot interact with each other. That is, Contoso R&D Net is isolated from Contoso Sales Net even though both are owned by Contoso. Contoso R&D Net contains three virtual subnets. Note that both the RDID and VSID are unique within a datacenter

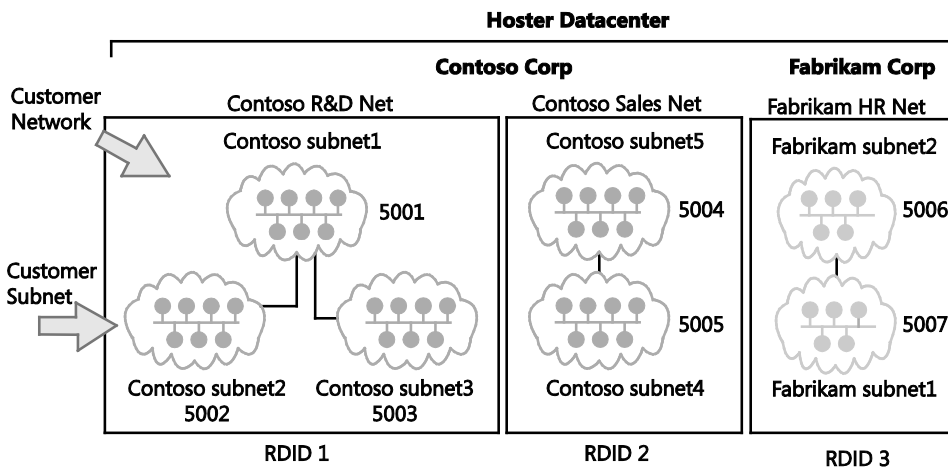


FIGURE 1-9 Example of RDIDs in a single datacenter

The VMs with VSID 5001 can have their packets routed or forwarded by network virtualization to VMs with VSID 5002 or VSID 5003. Before delivering the packet to the Hyper-V vSwitch, network virtualization will update the VSID of the incoming packet to the VSID of the destination VM. This will happen only if both VSIDs are in the same RDID. If the VSID that is associated with the packet does not match the VSID of the destination VM, the packet will be dropped. Therefore, virtual network adapters with RDID1 cannot send packets to virtual network adapters with RDID2.

In a physical network, a subnet is the L2 domain where computers (virtual and physical) can directly communicate with each other without having to be routed. In Windows, if you statically configure a network adapter, you can set a default gateway, which is the IP address to send all traffic that is going out of the particular subnet so that it can be routed appropriately. This is typically the router for your physical network. Hyper-V Network Virtualization uses a built-in router that is part of every host to form a distributed router for a virtual network. This means that every host, in particular the Hyper-V vSwitch, acts as the default gateway for all traffic that is going between virtual subnets that are part of the same VM network. In Windows Server 2012 and Windows Server 2012 R2, the address used as the default gateway is the lowest entry for the subnet (for example, it is the .1 address for a /24 subnet prefix). This address is reserved in each virtual subnet for the default gateway and cannot be used by VMs in the virtual subnet.

Hyper-V Network Virtualization acting as a distributed router provides a very efficient way to appropriately route all traffic inside a VM network because each host can directly route the traffic to the appropriate host without needing an intermediary. This is particularly true when two VMs in the same VM network but different virtual subnets are on the same physical host. As you will learn later, the packet never has to leave the physical host.

Most customer deployments will require communication from the Hyper-V Network Virtualization environment to resources that are not part of the Hyper-V Network Virtualization environment. Network virtualization gateways are required to allow communication between the two environments. Gateways can come in different physical form factors. They can be built on Windows Server 2012 R2, incorporated into a Top of Rack (TOR) switch or a load balancer, put into other existing network appliances, or can be a new standalone network appliance. As part of this POC configuration, you will deploy and configure a network virtualization gateway that is based on Windows Server 2012 R2. You'll learn more about this shortly.

Before you consider the gateway, it's important to understand how network packets actually flow within and between VM networks. You first need to understand two important terms:

- **Customer address (CA)** The CA is the IP address that the customer or tenant assigns based on the intranet infrastructure. This address lets the customer exchange network traffic with the VM as if it had not been moved to a public or private cloud. The CA is visible to the VM and reachable by the customer. Essentially, this is the IP address you would see if you ran an `ipconfig` command inside the VM.
- **Provider address (PA)** The PA is the IP address that the cloud service provider or the datacenter administrator assigns based on physical network infrastructure. The PA appears in the packets on the network that are exchanged with the server running Hyper-V that is hosting the VM. The PA is visible on the physical network, but not to the VM.

The CAs maintain the customer's network topology, which is virtualized and decoupled from the actual underlying physical network topology and addresses, as implemented by the PAs. The diagram in Figure 1-10 shows the conceptual relationship between VM CAs and network infrastructure PAs as a result of network virtualization.

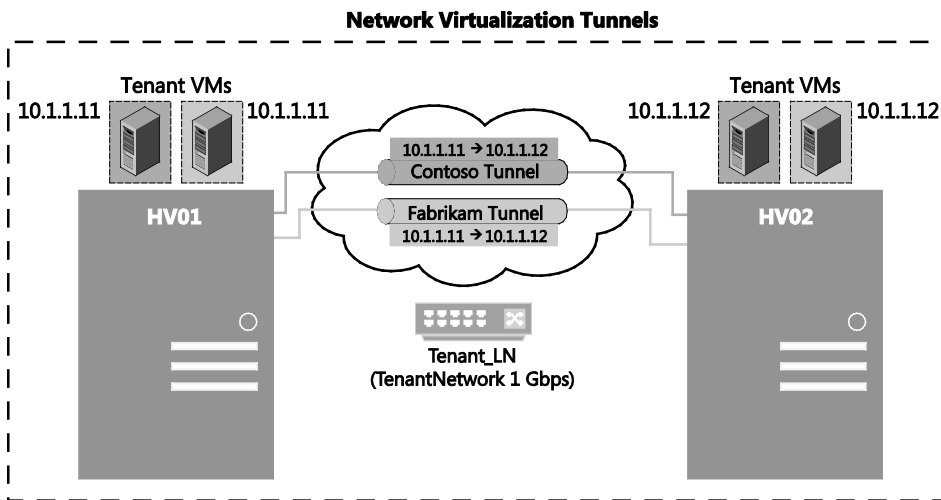


FIGURE 1-10 Packet delivery example using network virtualization

In Figure 1-10, customer VMs are sending data packets in the CA space. The packets traverse the physical network infrastructure through their own virtual networks, or tunnels. In Figure 1-10, you can think of the tunnels as envelopes that wrap around the Contoso and Fabrikam data packets. Shipping labels (PA addresses) determine how packets travel from the source host on the left to the destination host on the right. The key is how the hosts determine the shipping addresses (PAs) corresponding to the Contoso and the Fabrikam CAs, how the envelope is put around the packets, and how the destination hosts can unwrap the packets and deliver to the Contoso and Fabrikam destination VMs correctly. This simple analogy highlights the key aspects of network virtualization:

- Each VM CA is mapped to a physical host PA. Multiple CAs can be associated with the same PA.
- VMs send data packets in the CA spaces. These data packets are put into an envelope with a PA source and destination pair based on the mapping.
- The CA-PA mappings must allow the hosts to differentiate packets for different customer VMs.

As a result, the mechanism to virtualize the network is to virtualize the network addresses that the VMs use. Hyper-V Network Virtualization supports Network Virtualization for Generic Routing Encapsulation (NVGRE) as the mechanism to virtualize the IP address. NVGRE is part of the tunnel header. In NVGRE, the VM's packet is encapsulated inside another packet. The header

of this new packet has the appropriate source and destination PA IP addresses in addition to the VSID, which is stored in the Key field of the GRE header, as shown in Figure 1-11.

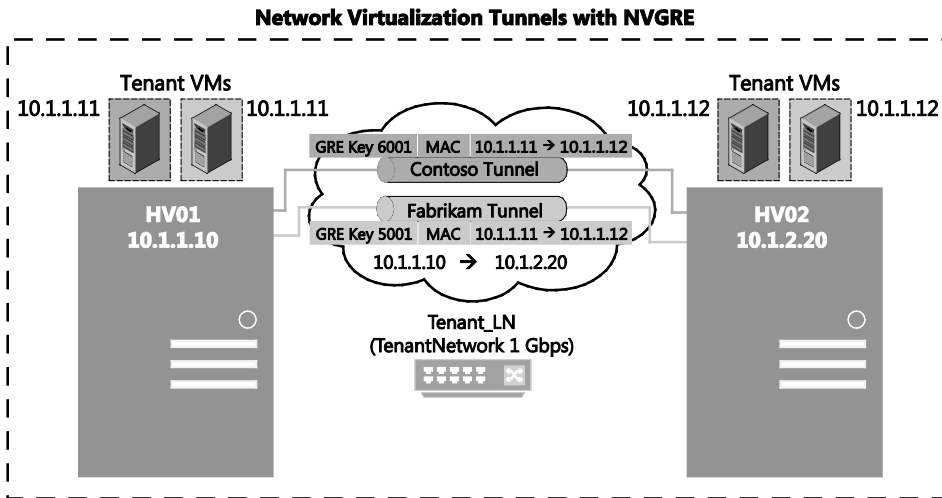


FIGURE 1-11 Example of NVGRE

The VSID allows hosts to identify the customer VM for any given packet, even though the PAs and the CAs on the packets may overlap. This allows all VMs on the same host to share a single PA, as shown in Figure 1-11.

Sharing the PA has a big impact on network scalability. The number of IP and MAC addresses that the network infrastructure needs to learn can be substantially reduced. For instance, if every end host has an average of 30 VMs, the number of IP and MAC addresses that the networking infrastructure needs to learn is reduced by a factor of 30. The embedded VSIDs in the packets also enable easy correlation of packets to the actual customers.

With Windows Server 2012 and later, Hyper-V Network Virtualization fully supports NVGRE out of the box; it does not require upgrading or purchasing new network hardware such as specialized network adapters, switches, or routers. This is because the NVGRE packet on the wire is a regular IP packet in the PA space, which is compatible with today's network infrastructure.

Windows Server 2012 made working with standards a high priority. With key industry partners (Arista, Broadcom, Dell, Emulex, Hewlett Packard, and Intel), Microsoft published a draft RFC that describes the use of GRE, which is an existing IETF standard, as an encapsulation protocol for network virtualization. As NVGRE-aware becomes commercially available, the benefits of NVGRE will become even greater.

Now that you understand how packets flow between VMs over the physical network, the final important piece to understand is how those packets can leave the virtualized networks entirely. For that, as mentioned earlier, you need a gateway.

As part of this POC configuration, you will deploy a total of four VM networks, each representing different departments within the environment. In addition, to more fully explore the different network virtualization capabilities, you will deploy a dedicated Hyper-V node together with two gateway VMs. These will collectively provide the gateway functionality for the aforementioned VM networks. You will deploy GW01 using the same bare-metal deployment capabilities you used for the storage and compute nodes. The configuration will resemble Figure 1-12.

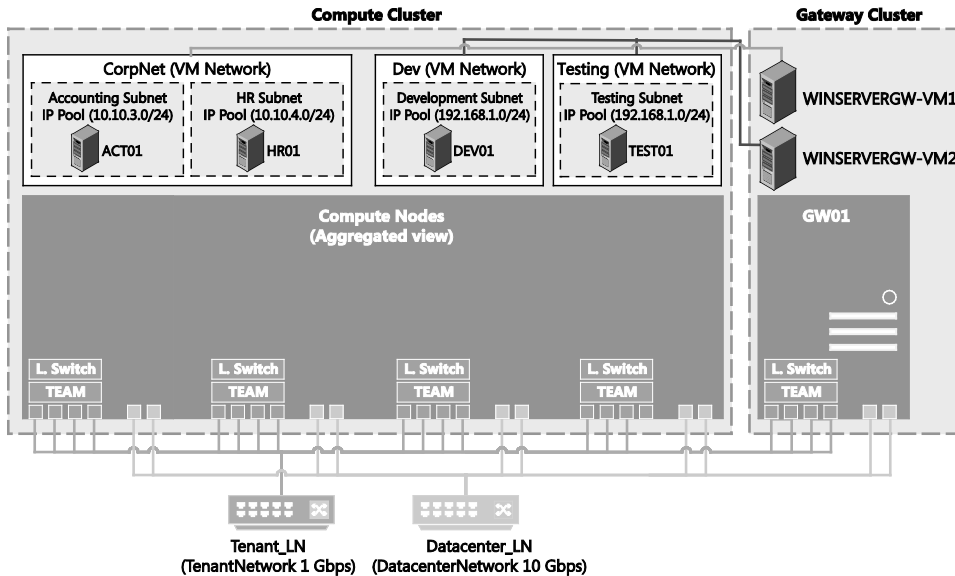


FIGURE 1-12 NVGRE gateway logical architecture

As shown in Figure 1-12, two Windows Server Gateway VMs are running on a single Hyper-V gateway node, GW01. You will dedicate this Hyper-V node to providing Windows Server Gateway functionality and place no other VMs on the node. In a production environment, it is a best practice that the VMs that will be providing the Windows Server Gateway services run on an underlying Hyper-V cluster, which, again, is dedicated to providing Windows Server Gateway services. To accelerate and streamline the deployment of the two Windows Server Gateway VMs, WINSERVERGW-VM1 and WINSERVERGW-VM2, you will use a predefined template that Microsoft provides. This will deploy two instances of the Windows Server Gateway; however, you will be deploying them with two different scenarios in mind. In a production environment, for high availability of network resources, you deploy the Windows Server Gateway with failover clustering by using two physical host servers running Hyper-V that are each also running a Windows Server Gateway VM that is configured as a gateway. The gateway VMs are then configured as a cluster to provide failover protection against network outages and hardware failure. The template Microsoft provides orchestrates these steps for you.

The first of the two scenarios that you will configure in this POC is aimed primarily at private cloud environments. Private cloud is a computing model that uses infrastructure dedicated to your organization. A private cloud shares many of the characteristics of public cloud computing, including resource pooling, self-service, elasticity, and metered services delivered in a standardized manner with the additional control and customization available from dedicated resources.

The fundamental difference between a private cloud and a public cloud is that a public cloud provides cloud resources to multiple organizations while the private cloud hosts resources for a single organization. However, a single organization may have multiple business units and divisions that can make a multitenant environment desirable. In these circumstances, the private cloud shares many of the security and isolation requirements of the public cloud.

For enterprises that deploy an on-premises private cloud, the Windows Server Gateway can act as a forwarding gateway and route traffic between virtual networks and the physical network. In this configuration, you will create two VM networks for two departments, accounting and HR. However, many of your key resources (such as Active Directory Domain Services and DNS) will be on your physical network (or at least, in a non-virtualized network). The Windows Server Gateway can route traffic between the virtual network and the physical network to provide employees working on the virtual network with all of the services they need. You will need to update your switches to ensure that servers on 10.10.0.0/24, 10.10.1.0/24, and 10.10.2.0/24 can reach 10.10.3.0/24 and 10.10.4.0/24. Otherwise, you should find that the workloads you deploy into the respective accounting and HR subnets behave exactly as they would if they had been placed within individual VLANs and routing had been configured between VLANs accordingly. The difference here is you configure this all with software and with very minimal switch reconfiguration.

The second scenario will use the other Windows Server Gateway VM, which this time, will be configured for network address translation (NAT). In this scenario, you have two other departments, development and testing, that need access to VMs, yet they need to remain away from CorpNet and the other management services. If you create two VM networks, each with an IP pool, your development and test teams will be able to deploy their own virtual workloads within isolated environments. Through the Windows Server Gateway NAT functionality, each group will still have external access to the Internet, for example, from within the VMs. WINSERVERGW-VM2 handles the translation from the 192.168.1.0/24 space inside the VMs and the 10.10.0.254 default gateway to allow connectivity to the outside world. Furthermore, inside the respective development and test VM networks, both use the 192.168.1.0/24 address space. No conflicts will occur because Hyper-V network virtualization will isolate each from the other, yet uniquely identify each by means of the unique IDs associated with the VM network and virtual subnet.

In a production environment, to ensure the highest levels of isolation, you should have a dedicated external or front-end network that would allow the development and testing VM networks to reach the Internet, rather than reaching it through the existing 10.10.0.254/24 address. But for testing and POC purposes, the simplified configuration will provide the same result.

By the end of Chapter 6, "Configuring network virtualization," you will have completed your deployment of Hyper-V with software-defined storage and networking. You'll have used System Center Virtual Machine Manager extensively to deploy, configure, and manage all elements of your fabric, and you will be ready to continue your evaluation of Windows Server 2012 R2, Hyper-V, and additional System Center components.

This page intentionally left blank

Deploying the management cluster

Chapter 1, “Design and planning,” familiarized you with the overall design of a proof-of-concept (POC) deployment and the aims of the solution. This chapter describes how to construct the infrastructure, starting with the key management tools for deploying and subsequently managing the solution.

Overview of the management cluster

For this POC configuration, infrastructure virtual machines (VMs) like the domain controller (DC), DNS server, Windows Deployment Services (WDS), Windows Server Update Services (WSUS), System Center Virtual Machine Manager (SCVMM), and other infrastructure services are deployed on top of a management cluster that consists of two Windows Server 2012 R2 Hyper-V nodes. In this infrastructure, each node in the management cluster runs on a Dell PowEredge R620 with 128 gigabytes (GB) of RAM, a RAID1 C:\ volume, and a RAID5 D:\ volume. From a networking standpoint, each node has four 1-Gbps network adapters and two 10-Gbps network adapters. The four 1-Gbps network adapters are for the TenantNetwork team, and the two 10-Gbps adapters are for the two redundant datacenter networks (DatacenterNetwork 1 and DatacenterNetwork 2). The network team configuration is discussed later in this chapter. During the initial phase of setup, the two management nodes (MGMT01 and MGMT02) remain unclustered until the storage cluster is deployed and available to the management cluster nodes as remote Server Message Block (SMB) storage. Initially, the management VMs use the local D:\ volume for VHDX file storage. After the storage cluster is configured and running, management VMs will be migrated to Scale-Out File Server storage for resiliency and increased performance.

Figure 2-1 shows the logical view of the management cluster, which consists of two Hyper-V servers named MGMT01 and MGMT02. The MGMT01 management node is on the left of the diagram, and the MGMT02 management node is on the right. The MGMT01 node consists of VMs DC01, VMM01, WDS, WSUS, and SQL01. DC01 is the first domain controller in the Contoso domain and also holds the DNS and Active Directory Domain Services roles. VMM01 runs System Center 2012 R2 Virtual Machine Manager. WDS runs Windows Deployment Services. WSUS runs Windows Server Update Services. SQL01 runs SQL Server 2012 SP1 and

houses the database for System Center 2012 R2 Virtual Machine Manager. MGMT02 is the second node of the management cluster, and it runs the VMs DC02, VMM02, and SQL02. DC02 is a replica domain controller for the Contoso domain. VMM02 runs as the failover VM in the highly available System Center 2012 R2 Virtual Machine Manager deployment. SQL02 runs SQL Server 2012 SP1 and hosts the AlwaysOn secondary replica of the System Center 2012 R2 Virtual Machine Manager database.

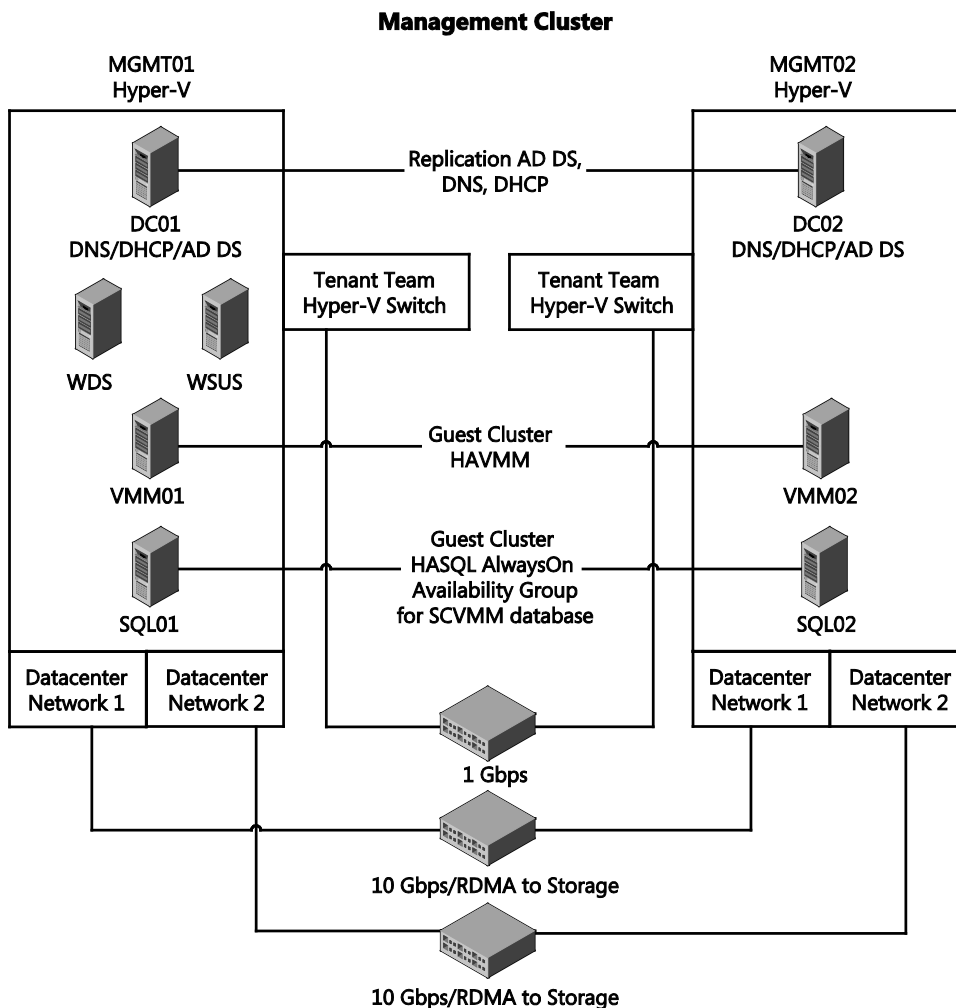


FIGURE 2-1 Logical view of the management cluster

The diagram illustrates the key concept that each management node has a NIC team that consists of four 1-Gbps network adapters. This NIC team is used exclusively for tenant traffic. The Tenant NIC team is connected to the Hyper-V tenant virtual switch on each management node to allow the VMs on the respective management node to communicate with network

resources outside of the Hyper-V environment. All of the network adapters in each NIC team are connected to a physical switch that provides network channel access to the management VMs on the node the switch is connected to. Additionally, the management nodes use two 10-Gbps network adapters that are set up to handle all non-tenant datacenter traffic, such as live migration and storage traffic. In summary, there are three physical networks: one for tenant traffic and two for datacenter traffic. These networks support the management infrastructure used in this POC.

The design shown in Figure 2-1 includes the configuration of redundant datacenter networks on all hosts and uses the 10-Gbps adapters for a variety of traffic types, including management, storage, and live migration. Windows Server, along with Data Center Bridging (DCB), handles the bandwidth allocation responsibilities of the converged network infrastructure. Again, each interface is capable of 10 Gbps and is redundant to layer 3. Routing is set up so that the network is self-healing in the event of a NIC, switch, or router failure. For an additional level of performance, the 10-Gbps adapters also support Remote Device Memory Access (RDMA) over Converged Ethernet (that is, the adapters are RoCE-capable). This capability can significantly boost performance for the key workloads SQL Server, Exchange Server, and other mission-critical applications and services.

Virtual Machine Manager

In this environment, Microsoft System Center 2012 R2 Virtual Machine Manager is the key application that runs on the management cluster. The management cluster, which is part of the compute fabric of the overall solution, provides management services for the environment with Virtual Machine Manager, Active Directory Domain Services, Domain Name Services (DNS), Windows Deployment Services, Windows Server Update Services (WSUS), and a variety of other infrastructure VMs needed to support the management of the environment. The management cluster is a collection of compute nodes that will form a single Hyper-V failover cluster. The key workloads on top of the management cluster, such as Virtual Machine Manager, provide necessary management functionality for the management cluster itself, the storage cluster, and compute cluster, as well as network devices. The management cluster utilizes storage from a storage cluster within the environment or may use resilient Direct Access Storage (DAS) with multiple domain controllers running on different management nodes.

This POC assumes that you will use Virtual Machine Manager to deploy your storage cluster. This means you need to deploy Virtual Machine Manager on local storage first because the storage cluster will not exist when you deploy Virtual Machine Manager.

For redundancy of the Active Directory Domain Services role, this POC uses two domain controllers, named DC01 and DC02, both running as VMs. Both are domain controllers for the contoso.com domain and run a forest- and domain-functional level of Windows Server 2012 R2. After you set up the domain controllers VMs, you will deploy additional VMs to host the Windows Deployment Services and WSUS roles. In this environment, you will deploy one VM for each of these roles. (If you need additional redundancy in your environment, you may deploy

multiple VMs for each role.) Next, you will deploy two SQL Server 2012 SP1 VMs in an AlwaysOn Availability Group to provide a resilient back end for the Virtual Machine Manager database. Then you will deploy Virtual Machine Manager in a pair of VMs, VMM01 and VMM02, in a highly available guest cluster configuration to ensure fault tolerance in the event of a physical host failure and VM-level failure. These VMs will be distributed across the two management nodes (MGMT01 and MGMT02) and stored, for the time being, on the D:\ volume of each node.

Service provider vs. enterprise

In a service provider infrastructure-as-a-service (IaaS) model, tenant workloads generally do not use the services of a management cluster. Usually, tenants deploy required infrastructure such as Active Directory Domain Services, DNS, and WSUS as part of the tenant workloads, and such workloads are isolated from other tenants by means of network virtualization or VLANs.

In an enterprise IaaS model, some or all of the tenant workloads may or may not use services from the management cluster, depending on whether an enterprise running an IaaS solution is the only tenant in the stamp.

Configuration walkthrough

This section walks through the steps for setting up the management cluster, from the initial hardware configuration and key considerations, to the VM creation and subsequent configuration.

This hardware infrastructure uses two identically configured Dell PowerEdge R620 servers for the management cluster. They include the following:

- 128 GB of RAM, dual Intel Xeon E5-2650 @ 2GHz, each with eight cores
- Four teamed 1-Gbps NICs, used for the TenantNetwork
- Two 10-Gbps RoCE-capable NICs, used for DatacenterNetwork 1 and DatacenterNetwork 2

Procedure 1: Rack and connect management cluster hosts

Use the following procedure to physically deploy the hardware into the environment. You can acquire new hardware or repurpose existing hardware for this procedure. If you already have hardware racked and cabled in your environment, ensure that the networking is configured as specified in this procedure.

The two management nodes are named MGMT01 and MGMT02. Perform the Procedure 1 steps on both MGMT01 and MGMT02. (Note that at this point, you are just plugging in the hardware. You'll configure IP addresses and NIC teams in subsequent procedures.)

1. Rack both physical servers that will become MGMT01 and MGMT02.
2. Connect power.

3. Connect the network. The management cluster hosts will use three separate networks:
 - DatacenterNetwork Port 1 (10 Gbps, not teamed)
 - DatacenterNetwork Port 2 (10 Gbps, not teamed)
 - TenantNetwork Team (teamed, consisting of four 1-Gbps adapters)
4. Plug all NICs on each node into physical network ports across separate switches. Note that these switches should be connected to one another to enable communication between the different networks.
 - Connect DatacenterNetwork Port 1 on each node to a switch that is RoCE-capable.
 - Connect DatacenterNetwork Port 2 on each node to another switch that is also RoCE-capable.
 - Connect half of the ports for TenantNetwork Team to one switch and half to another switch for redundancy.

Procedure 2: Install Windows Server on MGMT01 and MGMT02

The following procedure installs Windows Server 2012 R2 onto your bare-metal servers. Perform Procedure 2 on both MGMT01 and MGMT02. This procedure assumes that neither of your servers have an existing operating system or configuration installed locally. If your servers do have an existing configuration, it will be overwritten as part of this deployment. After completing this procedure, you can enable the Hyper-V role. It's important to note that this procedure shows the manual approach to deployment. If you have existing automated operating system deployment capabilities within your environment, you can skip this section. However, ensure that your systems have the latest firmware and optimal drive configuration.

1. Before you start, ensure that you have Windows Server 2012 R2 media. This media can be in the form of an ISO image, burned to a DVD, or expanded and transferred to a bootable USB stick. The installation of Windows Server 2012 R2 may vary slightly, depending on your media (retail, evaluation, volume licensed, and so on). You can download an evaluation version of Windows Server 2012 R2 from <http://www.microsoft.com/en-us/evalcenter/evaluate-windows-server-2012-r2>.
2. Power on your hosts. Ensure your management hosts have been configured to boot from your chosen media type. At the appropriate point in the boot process, enter the local RAID configuration utility and create a RAID1 drive, which will become your C:\ volume. Optionally, if you have additional disks, create a RAID5 drive, which will become your D:\ volume. If you do not have the appropriate number of disks for RAID5, choose a RAID level appropriate to your configuration. Exit the configuration utility.
3. When prompted, select the option to boot from your particular media.
4. When the Windows Setup dialog box appears, make the appropriate selections for language, time, and currency, as well as keyboard, based on your installation. Click Next.

5. In the next window, click Install Now.
6. On the Select The Operating System You Want To Install page, select either the Windows Server 2012 R2 Standard (Server with a GUI) or the Windows Server 2012 R2 Datacenter (Server with a GUI) edition, and then click Next. Choosing the full GUI will make local administration simpler and more straightforward. You can disable the GUI at a later point, if you prefer. This POC selects Datacenter edition for this deployment. However, both editions have feature parity, so if Standard is your preferred edition, all of the following steps will still apply.
7. On the License Terms page, click I Accept The License Terms, and then click Next.
8. For installation type, select Custom: Install Windows Only (Advanced). On the Where Do You Want To Install Windows page, select the drive that corresponds to your RAID1 disk, and click Next to start installation.

Installation may take a few minutes. At the end of the installation, after a restart, supply a password for the local administrator. You will use this password throughout this configuration.

9. After installation is complete and you have signed in as the local administrator, use Disk Management to configure the other local RAID array (if applicable). Format it as a simple volume with drive letter D:\ and make the partition type GPT. (Skip this step if your servers do not have additional disks for VM storage since you will place the management VMs on C:\ instead of D:\. Eventually, when configuration of the storage cluster is complete, most management VMs will use Scale-Out File Server storage.)
10. Update firmware as necessary to support NICs, array controllers, BMCs, and so on. Updating is important because if you are repurposing existing hardware for this POC, those firmware updates may unlock certain features and capabilities, as well as higher levels of performance and reliability.
11. Run Microsoft Update to update MGMT01 and MGMT02. This step uses Microsoft Update directly because it assumes you don't already have an internal update source and that you are deploying in an environment isolated from your production environment.
12. Acquire the latest Windows Server 2012 R2 network adapter drivers from your hardware vendor to support RoCE on the 10-Gbps network adapters.
13. Open Server Manager, click Local Server, and then click Computer Name in the Properties pane.
14. Click Change, replace the existing name with **MGMT01** or **MGMT02**, and click OK. When you are prompted, restart the server.

Procedure 3: Configure TenantNetwork teams

Perform Procedure 3 on both MGMT01 and MGMT02. The POC has four 1-Gbps NICs on each server and two 10-Gbps RoCE-capable NICs on each server. You will use Windows Server 2012

R2 NIC Teaming to combine the four 1-Gbps NICs into a single NIC team. Name the NIC team TenantNetwork Team.

As mentioned earlier, the TenantNetwork teams will combine four 1-Gbps NICs into a single team, which the VMs running on each management host will subsequently use to communicate with the rest of the network. This POC uses these four 1-Gbps NICs because the aggregated bandwidth will provide more than adequate performance. These NICs also provide redundancy against individual NIC failure. MGMT01 and MGMT02 will not use TenantNetwork Team for their own traffic. They will use the datacenter networks, DatacenterNetwork 1 and DatacenterNetwork 2, for that purpose.

1. Open Server Manager, and click Local Server on the left.
2. Under Properties, click one of your network adapters to open the Network Adapters window.
3. Right-click each network adapter and select Rename. Rename each network adapter as indicated in Table 2-1. A clear naming approach helps with identification and management later.

TABLE 2-1 Physical NIC and NIC team mapping

PHYSICAL NIC	PURPOSE	NAME
1-Gbps NIC #1	TenantNetwork Team	TenantNetwork Port 1
1-Gbps NIC #2	TenantNetwork Team	TenantNetwork Port 2
1-Gbps NIC #3	TenantNetwork Team	TenantNetwork Port 3
1-Gbps NIC #4	TenantNetwork Team	TenantNetwork Port 4
10-GB RoCE-capable NIC #1	DatacenterNetwork 1	DatacenterNetwork Port 1
10-GB RoCE-capable NIC #2	DatacenterNetwork 2	DatacenterNetwork Port 2

4. In Server Manager, under Properties, next to NIC Teaming, click the word Disabled to open the NIC Teaming window.
5. In the NIC Teaming window, next to Teams, click Tasks, click New Team, and select the check boxes next to your four TenantNetwork ports.
6. Name the team **TenantNetwork Team**, and then expand Additional Properties.
7. Under Teaming Mode, select Switch Independent. This mode uses algorithms that do not require the physical switch to participate in the teaming. Select the solution that best reflects your switches and environment. You can find more information about teaming in Windows Server 2012 R2 in the following whitepaper: <http://www.microsoft.com/en-us/download/details.aspx?id=40319>. For reference:
 - Static Teaming requires configuration on the switch and the computer to identify which links form the team.
 - The Link Aggregation Control Protocol (LACP) dynamically identifies links between the computer and a specific switch.

8. For the Traffic Distribution Algorithm, select Dynamic, which combines the best aspects of the Hyper-V Switch Port and Address Hashing modes into a single mode.
9. Under Standby Adapter, select None (all adapters Active). Your settings should reflect those shown in Figure 2-2. Confirm your settings, and click OK to create the team.
10. Repeat steps 1 through 9 on the other management node.

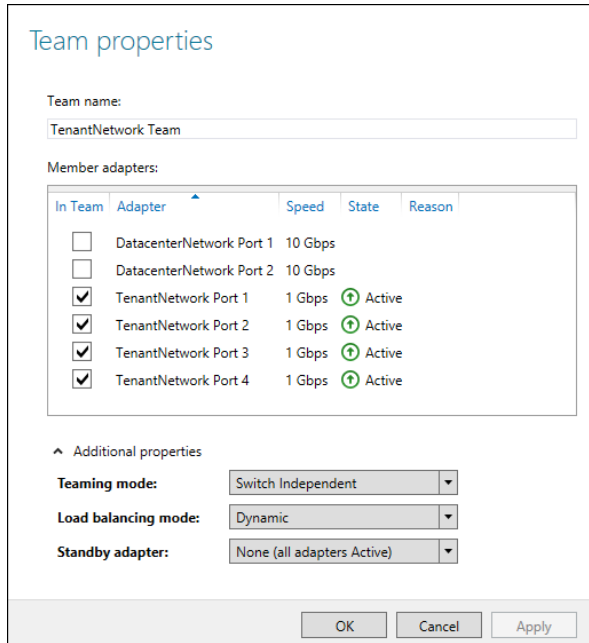


FIGURE 2-2 Tenant NIC team

11. Configure the IP addresses as shown in Table 2-2. Use the configuration data in Table 2-2 to configure the IP address information for the specified physical NIC.

TABLE 2-2 IP address assignments for NIC teams

PHYSICAL NIC	PURPOSE	NAME
DatacenterNetwork Port 1 on MGMT01	DatacenterNetwork 1	IP: 10.10.1.1 SM: 255.255.255.0 DG: 10.10.1.254
DatacenterNetwork Port 2 on MGMT01	DatacenterNetwork 2	IP: 10.10.2.1 SM: 255.255.255.0
DatacenterNetwork Port 1 on MGMT02	DatacenterNetwork 1	IP: 10.10.1.2 SM: 255.255.255.0 DG: 10.10.1.254
DatacenterNetwork Port 2 on MGMT02	DatacenterNetwork 2	IP: 10.10.2.2 SM: 255.255.255.0

Procedure 4: Enable the Hyper-V role on MGMT01 and MGMT02

With the networks configured, you can transform the existing Windows Server 2012 R2 servers into virtualization hosts by enabling the Hyper-V role. After the role is enabled, you can create the appropriate virtual switches to allow VM communication, and then you can create the management VMs that will control the infrastructure.

Follow these steps to enable the Hyper-V role using Server Manager.

1. On MGMT01, in Server Manager, click Manage, and then click Add Roles And Features.
2. On the Before You Begin page, click Next.
3. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
4. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select MGMT01, and then click Next.
5. On the Select Server Roles page, select Hyper-V. The Add Roles And Features dialog box appears. Click Add Features, and then click Next.
6. On the Create Virtual Switches page, ensure none of the boxes are selected, and then click Next.
7. Accept the defaults on the Virtual Machine Migration And Default Stores page, clicking Next on each page.
8. On the Confirm Installation Selections page, select Restart The Destination Server Automatically If Required, and then click Install. The server will restart automatically at the appropriate point.
9. After the server reboots, check Microsoft Update again to determine whether any additional updates now apply. Install updates and restart if necessary. Repeat all steps from this section on MGMT02.

The preceding steps describe how to use Server Manager to enable the Hyper-V role. Alternatively, you can use the Deployment Image Servicing and Management (DISM) cmdlets in Windows PowerShell to enable the role. To use Windows PowerShell, open Windows PowerShell as an administrator, and type:

```
Install-WindowsFeature -Name Hyper-V -IncludeManagementTools -Restart
```

Procedure 5: Configure Hyper-V virtual switches

The Hyper-V virtual switch is a software-based layer-2 network switch that includes programmatically managed and extensible capabilities to connect VMs to both virtual networks and the physical network. In addition, the Hyper-V virtual switch provides policy enforcement for security, isolation and service levels. In this environment, the Hyper-V virtual switches will enable the VMs on MGMT01 and MGMT02 to access the physical network and

communicate with each other. As shown in Figure 2-3, this configuration will bind a Hyper-V virtual switch to the freshly created TenantNetwork Team. This configuration means all VMs that attach to this new virtual switch will share 4 Gbps of bandwidth split across the redundant four physical network adapters.

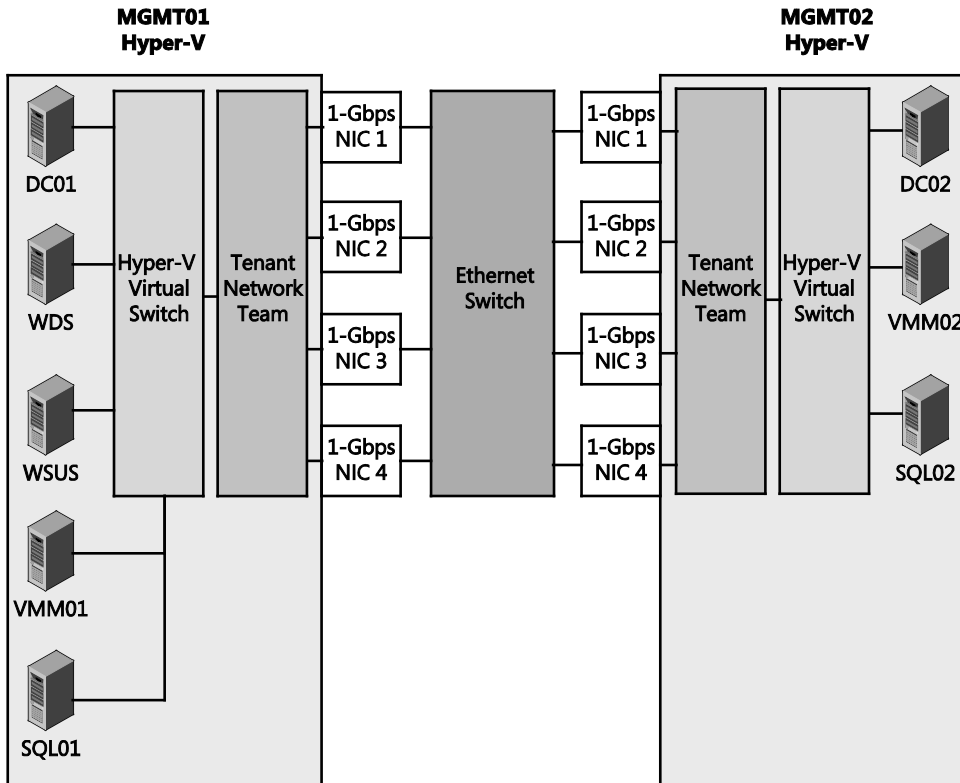


FIGURE 2-3 Logical mapping from the Hyper-V virtual switch to TenantNetwork Team

To specify this configuration, follow these steps:

1. On MGMT01, open Server Manager. On the Tools menu, click Hyper-V Manager.
2. Right-click MGMT01, and in the right pane, select Virtual Switch Manager.
3. Create a new external switch with the following settings, and then click OK:
 - **Name** TenantNetwork vSwitch
 - **Adapter** Microsoft Network Adapter Multiplexor Driver
 - **Allow Management OS to share this network adapter** Unchecked
4. In the Applying Networking Changes box, click Yes.
5. Repeat steps 1 through 4 on MGMT02.

You clear the Allow Management OS... check box so that the VMs have exclusive use of this particular virtual switch, and thus exclusive use of the underlying four 1-Gbps network

adapters. Two 10-Gbps adapters are already dedicated for host-related traffic such as storage, live migration, and management. These two dedicated adapters are more than enough. Therefore, the host does not need to share this virtual switch with the VMs. If you left the box checked, Hyper-V would create the virtual switch but would also create a virtual network adapter for the host operating system to allow it to communicate out of that particular virtual switch. Figure 2-3 shows how the logical network architecture of Hyper-V interfaces the operating system-level NIC team and physical NICs. Each VM communicates with the tenant Hyper-V virtual switch, which represents a synthetic layer-2 switch connected to TenantNetwork Team and allows each management VM to communicate over the physical network.

Procedure 6: Create folder structure for VMs and software

To simplify management and administration, follow these steps to create several folders on MGMT01 to hold key files associated with building the management cluster and the VMs.

1. On MGMT01, open File Explorer, and navigate to the D:\ volume. (If your management servers do not have a D:\ drive, use C:\ for folder paths.)
2. Right-click inside the D:\ volume window, select New, and then click Folder. Create the following folders on the D:\ volume.
 - \Exports
 - \Software
 - \Software\Server 2012 R2
 - \Software\Server 2012 R2\Extracted
 - \Software\VMM
 - \Software\VMM\Extracted
 - \Software\SQL
 - \Software\SQL\Extracted
 - \Software\ADK
 - \VMs
 - \VHDs
3. From the TechNet Evaluation Center, download the following key items and place them into the appropriate folders:
 - SQL Server 2012 SP1 Enterprise into \Software\SQL
<http://www.microsoft.com/en-us/evalcenter/evaluate-sql-server-2012-SP1>
 - System Center 2012 R2 Virtual Machine Manager into \Software\VMM
<http://www.microsoft.com/en-us/evalcenter/evaluate-system-center-2012-r2>
 - Windows 8.1 ADK into \Software\ADK
<http://go.microsoft.com/fwlink/?LinkId=309908>

- SQL Server 2012 SP1 utilities into \Software\SQL
<http://go.microsoft.com/fwlink/?LinkID=239648&clcid=0x409>
<http://go.microsoft.com/fwlink/?LinkID=239650&clcid=0x409>
 - Copy your existing Windows Server 2012 R2 ISO into \Software\Server 2012 R2; however, if you do not have an ISO, download the image from:
<http://www.microsoft.com/en-us/evalcenter/evaluate-windows-server-2012-r2>
4. After the image is downloaded, navigate to \Software\Server 2012 R2, right-click the ISO file, and select Mount. Copy the contents of the mounted ISO file to \Software\Server 2012 R2\Extracted.
 5. Navigate to \Software\SQL, right-click the ISO file, and select Mount. Copy the contents of the mounted ISO file to \Software\SQL\Extracted.
 6. Navigate to \Software\VMM. If you have downloaded Virtual Machine Manager from the TechNet Evaluation Center, you will have several System Center files. Double-click the SC2012_R2_SCVMM file, type **D:\Software\VMM\Extracted** as the extraction path when prompted, and then click Extract.
 7. Right-click the \Software folder, select Share With and Specific People. From the drop-down list, select Everyone, and then click Add. Click Share.
 8. After all extractions are finished, in File Explorer, click This PC, and then under Devices And Drives, right-click each DVD drive and select Eject to unmount the ISO files.

Procedure 7: Configure Hyper-V settings

With Hyper-V enabled, the virtual switches created, and the additional software downloaded into the newly created folder structure, you can finalize the Hyper-V settings before starting the deployment of the management VMs.

1. On MGMT01, in Hyper-V Manager, right-click MGMT01, and select Hyper-V Settings. Configure the default virtual hard disk location as D:\VHDs. (If your management servers do not have a D:\ volume, use C:\ for folder paths.)
2. Click Virtual Machines, configure the default VM location as D:\VMs, and then click OK.
3. Repeat steps 1 and 2 on MGMT02.
4. After both management hosts are configured, in Hyper-V Manager, in the top-right corner, click New, select Virtual Machine, use the following settings to complete the wizard, and then click Finish:
 - **Name** SYSPREP
 - **Location** D:\VMs
 - **Generation** 1
 - **Memory** 2048 MB (Dynamic Memory unchecked)
 - **Networking** TenantNetwork vSwitch

- **Virtual Hard Disk** Create a virtual hard disk with size 80 GB
 - **Installation Options** From an Image File, browse to the Windows Server 2012 R2 ISO from D:\Software\Server 2012 R2
5. After the VM is configured, right-click the SYSPREP VM, and select Start.
 6. Double-click the VM name to open a console window. When prompted to press any key to boot from CD/DVD, press any key.
 7. Perform a Windows Server 2012 R2 (Server with a GUI) installation. For detailed guidance, refer to Procedure 2, earlier in this chapter.
 8. After the installation starts, if possible, give the VM a temporary IP address and DNS information. Assign a default gateway to allow the use of Microsoft Update. Fully update the operating system, and reboot if necessary.
 9. Run Sysprep.exe from C:\Windows\System32\Sysprep, and select Generalize and Shutdown.
 10. Wait for the SYSPREP VM to shut down.
 11. Right-click the SYSPREP VM, and select Export. The SYSPREP VM is exported to D:\Export. A SYSPREP folder is created automatically.
 12. In Hyper-V Manager, select the SYSPREP VM, and then in the right pane, click Delete. You will use the exported SYSPREP VM to accelerate the creation of new VMs needed in the management cluster.

Procedure 8: Create management VMs

You can duplicate your new sysprepped image to quickly build the management infrastructure. Duplicating the image saves considerable time in contrast to attaching an ISO to each VM, installing the operating system, and then updating again. The deployment seems manual at this point. However, when the Virtual Machine Manager infrastructure is built, you can create and deploy templates centrally from the Virtual Machine Manager library instead of manually copying files on the network, attaching VHDX files to VMs, and so on.

To create a management VM, follow these steps on MGMT01:

1. On MGMT01, open File Explorer and navigate to D:\Exports.
2. Share D:\Exports as Exports with read and modify permissions to Everyone.
3. Copy the Sysprep.vhdx file inside D:\VHDs and rename the copy **DC01.vhdx**.
4. Copy the Sysprep.vhdx file inside D:\VHDs and rename the copy **WDS.vhdx**.
5. Copy the Sysprep.vhdx file inside D:\VHDs and rename the copy **WSUS.vhdx**.
6. Copy the Sysprep.vhdx file inside D:\VHDs and rename the copy **VMM01.vhdx**.
7. Copy the Sysprep.vhdx file inside D:\VHDs and rename the copy **SQL01.vhdx**.

Follow these steps on MGMT02:

1. Open File Explorer and navigate to \\MGMT01\Exports.
2. Right-click Sysprep.vhdx and copy it. Navigate to the local D:\ on MGMT02 and paste the file to D:\VHDs. Change the name to **DC02.vhdx**.
3. Paste the Sysprep.vhdx file again to D:\VHDs and change the name to **VMM02.vhdx**.
4. Paste the Sysprep.vhdx file again to D:\VHDs and change the name to **SQL02.vhdx**. You should now have three VHDX files in D:\VHDs on MGMT02.
5. Return to MGMT01, open Hyper-V Manager, and in the top-right corner, click New, and then click Virtual Machine.
6. Create a new virtual machine with the following settings:
 - **Name** DC01
 - **Location** D:\VMs
 - **Generation** 1
 - **Memory** 2048 MB (Dynamic Memory unchecked)
 - **Networking** TenantNetwork vSwitch
 - **Virtual Hard Disk** Use an existing VHD and browse to D:\VHDs\DC01.vhdx (If you have not set up a D:\ drive in your environment, then browse to C:\VHDs instead of D:\VHDs.)
7. Accept remaining defaults and create the virtual machine.
8. Create the remaining management VMs following Table 2-3. Observe the allocation of VMs across the two MGMT hosts.

TABLE 2-3 Management VMs with VHDX path, host, and memory setting

VM NAME	VHD FILE	HOST	MEMORY
DC01	D:\vhds\dc01.vhdx	MGMT01	2 GB (2048 MB)
DC02	D:\vhds\dc02.vhdx	MGMT02	2 GB (2048 MB)
VMM01	D:\vhds\vm01.vhdx	MGMT01	8 GB (8192 MB)
VMM02	D:\vhds\vm02.vhdx	MGMT02	8 GB (8192 MB)
SQL01	D:\vhds\sql01.vhdx	MGMT01	4 GB (4096 MB)
SQL02	D:\vhds\sql02.vhdx	MGMT02	4 GB (4096 MB)
WSUS	D:\vhds\wsus.vhdx	MGMT01	2 GB (2048 MB)
Windows Deployment Services	D:\vhds\wds.vhdx	MGMT01	2 GB (2048 MB)

Procedure 9: Configure Active Directory Domain Services on DC01

The following procedure configures the first management VM. To accelerate the process, use the sysprepped image you created and attached in Procedure 8. This method is considerably faster than mounting an ISO to each of the VMs and progressing through a full installation for each VM. With a few pieces of information, the VM will be up and running quickly, providing valuable functionality for the infrastructure, specifically, Active Directory Domain Services and DNS services. Follow this procedure on the DC01 VM running on MGMT01.

1. On MGMT01, open Hyper-V Manager and start DC01. Double-click DC01 to open a console to the VM.
2. After a few minutes, the out-of-box experience (OOBE) begins. Select the appropriate country or region, app language, and keyboard layout, and then click Next.
3. Accept the license terms, and then enter your local administrator password. (Use a standard password for all the local and domain accounts you set up during this POC.) Click Finish.
4. When the OOBE completes, log in to the VM with your local administrator credentials. Server Manager opens automatically. Click Local Server, and then click Computer Name in the Properties pane.
5. Click Change, replace the existing name with **DC01**, click OK, and then click OK in the Computer Name/Domain Changes dialog box. Reboot DC01.
6. When DC01 is back online, log in with your local administrator credentials.
7. In Server Manager, click Local Server, and then click the link next to Ethernet in the Properties pane.
8. Set the NIC information to the following:
 - **IP Address** 10.10.0.11
 - **Subnet Mask** 255.255.255.0
 - **Gateway** 10.10.0.254
 - **DNS** 127.0.0.1
9. In Server Manager, click Manage, and then click Add Roles And Features.
10. On the Before You Begin page, click Next.
11. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
12. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select DC01, and then click Next.
13. On the Select Server Roles page, select Active Directory Domain Services. On the Add Features That Are Required For Active Directory Domain Services page, click Add Features, and then click Next.

14. Click Next through the remaining pages, and then click Finish.
15. When the installation is complete, click the link Promote This Server To A Domain Controller To Finalize The Configuration Of This Domain Controller.
16. In the Active Directory Domain Services Configuration Wizard, select Add A New Forest. For the root domain name, type **contoso.com**, and then click Next.
17. Accept the default functional levels of Windows Server 2012 R2 for both the domain and forest. Type the same administrator password for the contoso.com domain that you used for the local administrator, and then click Next.
18. On the DNS Options page, you'll notice a warning at the top of the screen. Because this is the first DNS server in a fresh environment, you can safely ignore this warning. Click Next.
19. On the Additional Options page, click Next to accept the default.
20. On the Paths page, accept the defaults, and then click Next.
21. Review your settings on the Review Options page, and then click Next to initiate a prerequisite check. If all prerequisite checks pass successfully, click Install.
22. When prompted, restart DC01. When the server reboots, log in as contoso\administrator instead of the local administrator.

To configure DC01 with Windows PowerShell, with the same settings defined in the previous steps, run the following:

```
#Enable the ADDS Role:
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools

#Configure ADDS for the Contoso Domain
Install-ADDS-Forest -DomainName contoso.com -Confirm:$false
```

Procedure 10: Join management hosts to the Contoso domain

When the Contoso domain is established, to gain additional levels of control and simplified management, add the two management nodes, MGMT01 and MGMT02, to the Contoso domain.

1. Because DC01 is currently running on MGMT01, perform the first domain-join procedure on MGMT02.
2. Log on locally to MGMT02, and launch Server Manager.
3. In the left pane, click Local Server, and under Properties, next to Computer Name, click MGMT02.
4. In the System Properties window, click Change.
5. In the Computer Name/Domain Changes dialog box, select Domain, type **contoso.com**, and click OK.

6. When prompted for credentials to join the domain, use the contoso\administrator credentials that were established earlier, and click OK.
7. After the domain join is complete, accept the prompt to reboot the machine. Wait until MGMT02 is fully rebooted before moving on to the next step.
8. When MGMT02 is online and domain-joined successfully, log on to MGMT01 with the local administrator credentials.
9. Perform steps 2 through 7 on MGMT01.

After these steps are completed, both of your management nodes will be joined to the Contoso domain. When you reboot MGMT01, DC01's current running state will be saved and immediately resumed when you restart the node. This will not have an impact on the domain-join process in this POC environment.

Procedure 11: Configure DC02 as a secondary DC

The POC environment would likely have no issues with only a single DC. However, for redundancy and for load balancing of key services, you should add a secondary DC to the existing primary DC. In this environment, to add a further layer of redundancy, locate each DC on a separate management host to ensure that a single host being taken down won't impact the services Active Directory provides.

1. On MGMT01, open Hyper-V Manager, right-click Hyper-V Manager, select Connect To Server, and add MGMT02 to the Hyper-V Manager console on MGMT01.
2. Click MGMT02, and then start DC02. Double-click DC02 to open a console to the VM.
3. After a few moments, the OOBE begins. Select the appropriate country or region, app language, and keyboard layout, and then click Next.
4. Accept the license terms, enter your local administrator password, and then click Finish.
5. When the OOBE is complete, log in to the VM with your local administrator credentials. Server Manager opens automatically. Click Local Server, and then click Computer Name in the Properties pane.
6. Click Change, replace the existing name with **DC02**, click OK, and then click OK in the Computer Name/Domain Changes dialog box. Click Close, and then click Restart Now.
7. When DC02 is back online, log in with your local administrator credentials. In Server Manager, click Local Server, and then click the link next to Ethernet. In the Properties pane, enter the following details:
 - **IP Address** 10.10.0.12
 - **Subnet Mask** 255.255.255.0
 - **Default Gateway** 10.10.0.254
 - **DNS** 10.10.0.11
8. In Server Manager, click Manage, and then click Add Roles And Features.

9. On the Before You Begin page, click Next.
10. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
11. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select DC02, and then click Next.
12. On the Select Server Roles page, select Active Directory Domain Services. On the Add Features That Are Required For Active Directory Domain Services page, click Add Features, and then click Next.
13. Click Next through the remaining pages, and then click Finish.
14. When the installation is complete, click the Promote This Server To A Domain Controller link to finalize the configuration of this domain controller.
15. In the Active Directory Domain Services Configuration Wizard, select the option to add a domain controller to an existing domain. For the domain name, key **contoso.com**, and then click Next.
16. Next to Supply The Credentials To Perform This Operation, click Change, and then enter the contoso\administrator credentials.
17. Specify that DC02 is a DNS and Global Catalog Server to assist with name resolution and to hold copies of the different objects within the domain.
18. For the Directory Services Restore Mode (DSRM) password, enter the same credentials you used on DC01.
19. Click Next on all of the remaining pages, accepting the defaults. If all prerequisites pass, click Install.
20. After DC02 restarts, log on as contoso\administrator. It will take a short while for replication to complete. However, because the environment is currently very small, it shouldn't take long. You can proceed without waiting.

To configure DC02 with Windows PowerShell, with the same settings as defined in the previous steps, run the following:

```
#Enable the ADDS role:
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools

#Configure additional DC02 for the existing Contoso domain:
Install-ADSDomainController -InstallDNS -Credential (Get-Credential
contoso\administrator) -DomainName "contoso.com" -Confirm:$false
```

Procedure 12: Configure remaining management VMs

With the DCs running and providing authentication, name resolution, and other Active Directory services to the rest of the infrastructure, you can finalize the operating system

configuration on the remaining management VMs. These VMs will provide additional services in the environment.

Start with the System Center Virtual Machine Manager servers.

1. Open Hyper-V Manager on MGMT01.
2. Right-click VMM01, and select Start. Inside the VM, the OOB begins.
3. Double-click VMM01 to open a console, and set the local time zone and language settings.
4. Set the local administrator password to your standard administrator password. (Use a standard password for all the local and domain accounts you set up during this POC.)
5. For VMM01, open the network settings, set the following, and then click OK:
 - **IP Address** 10.10.0.13
 - **Subnet Mask** 255.255.255.0
 - **Default Gateway** 10.10.0.254
 - **Primary DNS** 10.10.0.11
 - **Secondary DNS** 10.10.0.12
6. Open Server Manager, click Local Server, and then click Computer Name in the Properties pane.
7. Click Change, replace the existing name with **VMM01**, and then select Domain.
8. Enter **contoso.com** for the domain name, and click OK.
9. When prompted for credentials, use contoso\administrator with your standard password.
10. Click OK, and then click Close. When prompted to restart, click Restart Now.

VMM01 restarts, and you can continue to configure the remaining VMs. Repeat steps 2 through 10 above on the remaining management VMs. Table 2-4 provides the details. Note that for each of the management VMs, the subnet mask is 255.255.255.0, the default gateway is 10.10.0.254, the primary DNS is 10.10.0.11, and the secondary DNS is 10.10.0.12.

TABLE 2-4 Management VMs and IP addresses

NAME	HOST	IP ADDRESS
VMM02	MGMT02	10.10.0.14
SQL01	MGMT01	10.10.0.15
SQL02	MGMT02	10.10.0.16
WSUS	MGMT01	10.10.0.17
WDS	MGMT01	10.10.0.18

Procedure 13: Install and configure WSUS on the WSUS VM

WSUS plays an important part in this environment. It helps to centralize the collection, administration, and deployment of updates that will be deployed to the key workloads, both physical and virtual, within this infrastructure. This procedure enables the WSUS role on the WSUS VM and configures an approval rule that will automatically approve certain updates. This will streamline the deployment in the POC.

1. Log on to your WSUS VM as contoso\administrator with your established password.
2. In Server Manager, click Manage, and then click Add Roles And Features.
3. On the Before You Begin page, click Next.
4. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
5. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select WSUS, and then click Next.
6. On the Select Server Roles page, select Windows Server Update Services. An Add Roles And Features dialog box appears. Click Add Features, and then click Next.
7. On the Select Features page, accept the default selections, and then click Next.
8. On the Windows Server Update Services page, click Next.
9. On the Select Role Services page, accept the default selections, and then click Next.
10. On the Content Location Selection page, type a valid location to store the updates. Use C:\WSUS, but note that in a production environment, you may locate this repository on a separate drive. Click Next.
11. The Web Server Role (IIS) page opens. Review the information, and then click Next. On the Select The Role Services To Install For Web Server (IIS) page, accept the defaults, and then click Next.
12. On the Confirm Installation Selections page, review the selected options, and then click Install. The WSUS Installation Wizard runs. This might take several minutes to complete.
13. After WSUS installation is complete, in the summary window on the Installation Progress page, click Launch Post-Installation Tasks. The text changes to Please Wait While Your Server Is Configured. When the task is finished, the text changes to Configuration Successfully Completed. Click Close.
14. In Server Manager, verify whether a required restart notification appears. This can vary according to the installed server role. If the notification appears, restart the server to complete the installation.

To configure WSUS with Windows PowerShell, run the following:

```
#Enable the WSUS Role:
Install-WindowsFeature -Name UpdateServices -IncludeManagementTools

#Post-installation, specify where to store downloaded updates:
cd 'C:\Program Files\Update Services\Tools'.\WsusUtil.exe PostInstall
CONTENT_DIR=C:\WSUS
```

With the WSUS role installed and its initial configuration complete, you can add configuration specific to this environment and the workloads it will include. Before you proceed, WSUS must be able to access the Internet. If a corporate firewall is between WSUS and the Internet, you might have to configure that firewall to ensure WSUS can obtain updates. To obtain updates from Microsoft Update, the WSUS server uses port 443 for the HTTPS protocol. Although most corporate firewalls allow this type of traffic, some companies restrict servers from accessing the Internet due to company security policies. For more detailed steps on configuring the network on the WSUS VM, see <http://technet.microsoft.com/en-us/library/hh852346.aspx>.

1. On WSUS, in Server Manager, click Tools, and then click Windows Server Update Services.
2. If the Complete WSUS Installation dialog box appears, click Run. In the Complete WSUS Installation dialog box, click Close when the installation successfully finishes.
3. The Windows Server Update Services Wizard opens. On the Before You Begin page, review the information, and then click Next.
4. Read the instructions on the Join The Microsoft Update Improvement Program page and evaluate whether you want to participate. To participate in the program, retain the default selection. Otherwise, clear the check box, and then click Next.
5. On the Choose Upstream Server page, select Synchronize The Updates With Microsoft Update, and then click Next.
6. On the Specify Proxy Server page, if you are using a proxy in your environment, select the Use A Proxy Server When Synchronizing check box, enter the details, and then click Next. If you are not using a proxy, just click Next.
7. On the Connect To Upstream Server page, click Start Connecting. When it connects, click Next.
8. On the Choose Languages page, select Download Updates Only In These Languages, and then select the languages you want updates for. For this configuration, choose only English, and then click Next.
9. On the Choose Products page, select the following:
 - SQL Server 2012 SP1
 - System Center 2012 R2 Virtual Machine Manager
 - Windows Server 2012 R2

10. On the Choose Classifications page, for simplicity with this deployment, select all classifications, and then click Next.
11. On the Set Sync Schedule page, select the options most relevant for your environment. For this POC, select Synchronize Automatically, and leave the remaining defaults. Click Next.
12. On the Finished page, select the Begin Initial Synchronization option to start the synchronization. This may take some time, depending on your Internet connection speed and the classifications you chose.

Procedure 14: Configure WSUS GPO and auto-approvals

With WSUS configured and updates being downloaded from Microsoft Update, you need to ensure that the new WSUS configuration is the centralized source for updates for all of the key servers, both physical and virtual, within this POC environment. To streamline the update process, configure a Group Policy Object (GPO) with the relevant WSUS settings. You need to create a new GPO that contains only WSUS settings. You can link this GPO to an Active Directory container that is appropriate to the environment. In a small POC like this, it is suitable to link the GPO to the domain GPO. However, in a more complex environment, you may want to restrict this linkage to only specific organizational units.

1. Log on to DC01 with the contoso\administrator credentials.
2. Open Server Manager, click Tools, and then click Group Policy Management.
3. Right-click contoso.com, and select Create A GPO In This Domain And Link It Here.
4. Enter the name **WSUS Settings**, and click OK.
5. Right-click the newly created WSUS Settings, and select Edit.
6. In the Group Policy Management Console (GPMC), expand Computer Configuration, expand Policies, expand Administrative Templates, expand Windows Components, and then click Windows Update.
7. In the details pane, double-click Configure Automatic Updates. The Configure Automatic Updates policy opens.
8. Click Enabled, and then select Auto Download And Notify For Install. This option automatically begins downloading updates and then notifies a logged-on administrative user before installing the updates.
9. Click OK to close the Configure Automatic Updates policy and return to the Windows Update details pane.
10. In the Windows Update details pane, double-click Specify Intranet Microsoft Update Service Location.
11. Click Enabled, set the Intranet Update Service for Detecting Updates and Set the Intranet Statistics Server text boxes. Type **http://wsus.contoso.com:8530**, and then click OK.

It may take a short time for Group Policy to apply the new policy settings to the other member servers in the domain. By default, Group Policy updates in the background every 90 minutes with a random offset of 0 to 30 minutes. If you want to update Group Policy sooner, you can open a command prompt window on the management VMs and the hosts, and type **gpupdate /force**.

With the WSUS-specific GPO configured and linked, the final step is to create an automatic approval rule in WSUS. In this POC, this step streamlines the approval of updates that have been downloaded from Microsoft Update for the desired products and classifications. In a production environment, you would substitute this automatic approval rule with one that is more tailored to your specific needs. This would, for example, involve approval for a subset of servers that undergo a full post-deployment test before approval for deployment to a broader set of servers.

1. On WSUS, in Server Manager, click Tools, and then click Windows Server Update Services.
2. In the WSUS Administration Console, under Update Services, expand the WSUS server, and then click Options. The Options window opens.
3. In the Options window, click Automatic Approvals. The Automatic Approvals dialog box opens.
4. In Update Rules, select the Default Automatic Approval Rule, and click Edit.
5. In Step 1: Select Properties, clear all boxes.
6. In Step 2: Edit The Properties, click the link, and select All Computers and Unassigned Computers. Click OK.
7. Click Run Rule. This may take a few moments. The window displays how many updates have been approved.
8. Click OK to close the Automatic Approvals dialog box.

The automatic approval rule simplifies updating the POC infrastructure and ensures that Windows Server, SQL Server, and System Center Virtual Machine Manager stay up to date.

Procedure 15: Install Windows Deployment Services

Windows Deployment Services allows you to deploy Windows Server instances and Windows client instances to target devices across the network. You can use Windows Deployment Services as a standalone method to accelerate deployment of a Windows operating system. But on its own, Windows Deployment Services typically requires manual input at the device side to enter relevant details, as was the case earlier when you installed Windows Server from the DVD. Windows Deployment Services, however, also serves as the deployment engine for more functional management tools such as Virtual Machine Manager, which uses Windows Deployment Services to centrally deploy new Hyper-V hosts over the network. Procedure 15 simply configures Windows Deployment Services so that it is ready to use with Virtual Machine Manager later.

1. Log on to your Windows Deployment Services VM with the contoso\administrator credentials.
2. In Server Manager, click Manage, and then click Add Roles And Features.
3. On the Before You Begin page, click Next.
4. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
5. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select WDS, and then click Next.
6. On the Select Server Roles page, select Windows Deployment Services. Click Add Features, and then click Next.
7. On the remaining pages of the wizard, click Next until you reach the confirmation page. Click Install.
8. When installation is complete, in Server Manager, click Tools, and then click Windows Deployment Services.
9. In the Windows Deployment Services window, expand Servers, right-click WDS, and select Configure Server.
10. In the Windows Deployment Services Configuration Wizard, select Integrated With Active Directory, and then click Next.
11. For this POC, leave the default C:\RemotelInstall; however, you can use an alternative if you prefer. Click Next, and then click Yes to accept the system volume warning.
12. For client selection, select Do Not Respond To Any Client Computers, and click Next.
13. On the Operation Complete page, clear the check box next to Add Images To The Server Now. (Virtual Machine Manager will handle this later.)

To configure Windows Deployment Services with Windows PowerShell, run the following:

```
#Enable the Windows Deployment Services role:  
Install-WindowsFeature -Name WDS -IncludeManagementTools
```

Procedure 16: Create administrative service accounts

Procedure 16 quickly sets up some credentials for the accounts that will be used to run the services for Virtual Machine Manager and SQL Server 2012 SP1. Although it is possible to use the contoso\administrator account for both of these accounts, a best practice is to use delegated accounts within the domain with enough privilege to perform the desired functionality.

1. Log on to your DC01 VM with the contoso\administrator credentials.
2. In Server Manager, click Tools, and then click Active Directory Users And Computers.
3. Expand contoso.com, and click the Users Organizational Unit (OU).

4. Right-click Users OU, select New, and then select User. Enter the following information, and then click Next:
 - **First Name** SCVMM_svc
 - **User Logon Name** SCVMM_svc
5. For the password, use your established credentials that you have been using throughout the setup. If you plan to keep the POC environment available for an extended period of time, check the box Password Never Expires.
6. Clear the User Must Change Password At Next Logon check box, click Next, and then click Finish.
7. Repeat steps 4 through 6 with SQL_svc as the account.
8. Log on to your VMM01 VM with the contoso\administrator credentials.
9. In Server Manager, click Tools, and then click Computer Management.
10. Expand Local Users And Groups, click Groups, and then double-click Administrators.
11. In the Administrators Properties window, click Add, select SCVMM_svc, and then click OK. Click OK again to close the Administrators Properties window. Close Computer Management.
12. Repeat steps 8 through 11 on VMM02.
13. Repeat steps 8 through 11 on SQL01 and SQL02 using the SQL_svc account.

Procedure 17: Add data and log disks to SQL01 and SQL02

To follow best practices for SQL Server deployment, use the Hyper-V Manager console to add D:\ and E:\ volumes to both SQL Server VMs. After adding the disks in Hyper-V Manager, initialize and format them in Disk Management, within the guest OS on the respective VMs.

1. Log on to your MGMT01 host with the contoso\administrator credentials.
2. Open Hyper-V Manager, right-click SQL01, and select Settings.
3. In the left pane, click SCSI Controller, and in the central pane, click Hard Disk.
4. Click Add. Select Virtual Hard Disk, and click New.
5. On the Before You Begin page, click Next.
6. On the Choose Disk Type page, select Dynamically Expanding. This creates a virtual hard disk with a physical size that represents the size of the virtual disk's contents. For a production environment, you can optionally convert this dynamically expanding disk to fixed size later. Click Next.
7. On the Specify Name And Location page, name the disk **SQL01_Ddisk.vhdx**, and ensure that D:\VHDs is the target folder. Click Next.
8. On the Configure Disk page, select the Create A New Blank Virtual Hard Disk option. Enter **40** in the box to create a 40-GB VHDX file. This can be expanded later without VM downtime. Click Next.

9. On the Completing The New Virtual Hard Disk Wizard page, review the summary, and click Finish.
10. In the Settings For SQL01 On MGMT01 window, click Apply to attach the new VHDX.
11. Repeat steps 3 through 10, using a disk name of SQL01_Edisk.vhdx. After completing these steps, click OK to close the SQL01 VM settings window.
12. In Hyper-V Manager, click MGMT02, click SQL02, then select Settings.
13. Repeat steps 3 through 11 to create the D:\ and E:\ volumes for SQL02 on MGMT02.
14. In Hyper-V Manager, log on to your SQL01 VM.
15. To open Disk Management, right-click Start, and select Disk Management.
16. Right-click the new disks, one by one, and select Online.
17. Right-click the first of your two new disks, and select Initialize.
18. Leave both new disks selected, select GPT for the partition type, and click OK.
19. Right-click the first new disk, and select New Simple Volume.
20. When the New Simple Volume Wizard opens, click Next.
21. On the Specify Volume Size page, accept the default, and click Next.
22. On the Assign Drive Letter Or Path page, select the Assign The Following Drive Letter option, and select D from the drop-down list. Click Next.
23. On the Format Partition page, select the Format This Volume With The Following Settings option.
24. Use the drop-down list to change the allocation unit size to 64 kilobytes. Label the volume **SQL_Data**, and click Next.
25. On the Completing The New Simple Volume Wizard page, review your selections, and click Finish.
26. Repeat steps 19 through 25 for the second drive on SQL01, this time using drive letter E and SQL_Logs as the label name.
27. Repeat steps 14 through 26 for the two new drives on SQL02, currently hosted on MGMT02.

Procedure 18: Add a second virtual network adapter to SQL01 and SQL02

To prevent a single point of failure with the failover cluster, add a second NIC to each cluster node. This procedure uses address range 192.168.1.0/24, but you can use any address range that works for your POC.

1. Log on to your MGMT01 host with the contoso\administrator credentials.
2. Open Hyper-V Manager.
3. Register MGMT02 in the console if it's not already registered.

4. Connect to your SQL01 VM and shut it down.
5. Click MGMT02 in the console.
6. Connect to your SQL02 VM and shut it down.
7. Right-click SQL01, and select Settings.
8. In the left pane, click Add Hardware, and select Network Adapter.
9. For the virtual switch, select the tenant virtual switch, and click OK.
10. Repeat steps 4 through 9 for SQL02.
11. Start both SQL01 and SQL02.
12. Log on to SQL01 using your contoso\administrator credentials.
13. From Server Manager, open the network adapter settings for any adapter. You will see two adapters. Rename the first adapter **TenantNetwork** and rename the second adapter **ClusterNetwork**.
14. Set the following IP address properties for the ClusterNetwork adapter:
 - **IP Address** 192.168.1.1
 - **Subnet Mask** 255.255.255.0
 - **DNS** Leave blank
 - **Default Gateway** Leave blank
15. Save the settings.
16. Log on to SQL02 using your contoso\administrator credentials.
17. From Server Manager, open the network adapter settings for any adapter. You will see two adapters. Rename the first adapter **TenantNetwork** and rename the second adapter **ClusterNetwork**.
18. Set the following IP address properties for the ClusterNetwork adapter:
 - **IP Address** 192.168.1.2
 - **Subnet Mask** 255.255.255.0
 - **DNS** Leave blank
 - **Default Gateway** Leave blank
19. Save the settings.

Procedure 19: Create the guest cluster within SQL01 and SQL02

To ensure high availability for the Virtual Machine Manager database, you can configure a Windows Server failover cluster on SQL01 and SQL02. You will then use it to configure a SQL Server AlwaysOn Availability Group. Although you will be clustering the two SQL Server VMs, you will not be setting SQL Server 2012 SP1 AlwaysOn in a Failover Cluster Instance (FCI) configuration because, at this time, shared storage is not yet configured as is required for

installing a SQL Server failover cluster. Instead, you will install SQL Server in a standalone configuration on both SQL01 and SQL02. After Virtual Machine Manager is installed, you will configure an AlwaysOn Availability Group for the Virtual Machine Manager database. This is a requirement because the two SQL Server servers are members of the same failover cluster. Therefore, you'll achieve high availability and automatic failover at the database level for the Virtual Machine Manager database. The following steps walk you through creating the failover cluster with SQL01 and SQL02 as members.

1. Log on to your SQL01 VM with the contoso\administrator credentials.
2. In Server Manager, click Manage, and then click Add Roles And Features.
3. On the Before You Begin page, click Next.
4. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
5. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select SQL01, and then click Next.
6. On the Select Server Roles page, click Next.
7. On the Select Features page, click Failover Clustering, and when prompted, on the Add Features That Are Required For Failover Clustering page, click Add Features, and then click Next.
8. On the Confirm Installation Selections page, click Install.
9. Repeat steps 1 through 8 on SQL02. When installations are complete on both servers, proceed.
10. On SQL01, open Server Manager, click Tools, and select Failover Cluster Manager.
11. Right-click Failover Cluster Manager, and select Create Cluster.
12. On the Before You Begin page, click Next.
13. On the Select Servers page, add SQL01 and SQL02, and click Next.
14. On the Testing Options page, select the Run All Tests (Recommended) option, and click Next.
15. On the Confirmation page, click Next. This validates the proposed cluster configuration against a predetermined set of tests and presents a report at the end of validation. A warning that shared storage is not yet configured appears, but the validation tests will pass. Optionally, you can view the report for more detail.
16. The Create Cluster Wizard opens. On the Before You Begin page, click Next.
17. On the Access Point For Administering The Cluster page, enter **SQLCluster** for the name, assign IP address **10.10.0.25** (from the tenant range), and then click Next.
18. On the confirmation page, review your selections, and click Next to start creating the cluster. This will take a few moments to complete.

To install Failover Clustering with Windows PowerShell, run the following:

```
#Enable Failover Clustering:  
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools
```

A new failover cluster named SQLCluster.contoso.com should now be formed with two nodes: SQL01 and SQL02. Having the two SQL Server nodes as members of the same failover cluster is a requirement for enabling AlwaysOn availability at the service level and also for creating an AlwaysOn Availability Group.

Procedure 20: Install SQL Server on SQL01 and SQL02

Procedure 20 describes how to install SQL Server 2012 SP1 Enterprise Edition on SQL01 and SQL02 to support the installation of Virtual Machine Manager. After SQL Server is installed on its own set of VMs, you can configure an AlwaysOn Availability Group for the Virtual Machine Manager database. The benefit of using two SQL Server servers to host the Virtual Machine Manager database is to provide high availability at the database level.

1. Log on to your SQL01 VM with the contoso\administrator credentials.
2. In Server Manager, click Manage, and then click Add Roles And Features.
3. On the Before You Begin page, click Next.
4. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option, and click Next.
5. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select SQL01, and then click Next.
6. On the Select Server Roles page, click Next.
7. On the Select Features page, expand .NET Framework 3.5 Features, select .NET Framework 3.5 (Includes .NET 2.0 and 3.0), and then click Next.
8. On the Confirm Installation Selections page, click Specify An Alternate Source Path.
9. On the Specify Alternate Source Path page, enter **\\MGMT01\Software\Server 2012 R2\Extracted\Sources\SxS**, and click OK.
10. On the Confirm Installation Selections page, click Install.
11. Run Windows Update to check for any .NET framework-specific updates.
12. Repeat steps 1 through 11 on SQL02.
13. On SQL01, open File Explorer, navigate to **\\MGMT01\Software\SQL\Extracted**, and then click Setup.
14. Click Installation in the left pane of the Setup page, and select New SQL Server Stand-Alone Installation Or Add Features To An Existing Installation.
15. Enter a license key or specify Free Evaluation (180 days).

16. On the License Terms page, review the license agreement. If you agree, select the I Accept The License Terms check box, and then click Next.
17. In the Global Rules window, the setup procedure automatically advances to the Product Updates window if there are no rule errors. You may receive a warning indicating that Windows Firewall is running and that you should verify that SQL Server ports are open in the firewall. For this installation, install a default instance of SQL Server, which means open inbound port 1433. For more information on opening a SQL Server firewall port for TCP and dynamic ports, see <http://msdn.microsoft.com/en-us/library/ms175043.aspx/>.
18. The Microsoft Update page appears if the Microsoft Update check box in Control Panel\All Control Panel Items\Windows Update\Change settings is not checked. Putting a check in the box on the Microsoft Update page will change the computer settings to include the latest updates when you scan for Windows Update.
19. On the Product Updates page, the latest available SQL Server product updates are displayed. If no product updates are discovered, SQL Server Setup does not display this page and auto-advances to the Install Setup Files page.
20. The Install Setup Files page shows the progress of downloading, extracting, and installing the setup files. If an update for SQL Server Setup is found and is specified to be included, that update is also installed.
21. On the Setup Role page, select SQL Server Feature Installation, and then click Next.
22. On the Feature Selection page, select the features for your installation. Select Database Engine Services and Management Tools - Complete. Accept the default installation paths, and click Next.
23. On the Instance Configuration page, specify the default instance, MSSQLSERVER.
24. On the Server Configuration - Service Accounts page, specify login accounts for SQL Server services. Set the SQL Server Agent and the SQL Server Database Engine to start up as Contoso\SQL_svc, and use the password you set when you created the service accounts. Leave the default selections for other service-related settings.
25. On the Database Engine Configuration page, click Add Current User. Set the data root directory to D:\, amend the User Database Log Directory and Temp DB Log Directory settings to E:\ instead of D:\, and then click Next.
26. On the Ready To Install page, review the settings, and then click Install.
27. After these steps are completed, run Windows Update to ensure you have the latest updates for your SQL Server installation, and restart if necessary.
28. Repeat steps 13 through 27 on SQL02.

To install the .NET Framework 3.5 features with Windows PowerShell, run the following:

```
#Enable the Windows Deployment Services Role:
Install-WindowsFeature -Name NET-Framework-Core -source
    "\\MGMT01\Software\Server 2012R2\Extracted\Sources\Sxs\"
```

Procedure 21: Enable AlwaysOn high availability for SQL01 and SQL02

As mentioned earlier, the use of SQL Server AlwaysOn Availability Groups will increase redundancy and native protection for the key Virtual Machine Manager database. In this configuration, because no shared storage is currently available, this procedure will use the Availability Replica functionality that exists within Availability Groups. This will enable you to define a primary replica for the Virtual Machine Manager database on the primary SQL Server node, SQL01, and create a corresponding secondary replica on the other SQL Server node, SQL02. This procedure will simply enable the AlwaysOn Availability Groups functionality and will continue the process after you've created the Virtual Machine Manager database in a later step. You can read more information on Availability Groups on MSDN at <http://msdn.microsoft.com/en-us/library/ff877884.aspx>.

1. Log on to your SQL01 VM with the contoso\administrator credentials.
2. Open SQL Server Configuration Manager.
3. In the left pane, click SQL Server Services.
4. In the right pane, right-click SQL Server (MSSQLSERVER), and select Properties.
5. Click the AlwaysOn High Availability tab.
6. Verify that the Windows Failover Cluster name contains the name of the local failover cluster. If it is blank, this server instance currently does not support AlwaysOn Availability Groups: The local computer is not a cluster node, the WSFC cluster has been shut down, or this edition of SQL Server 2012 SP1 does not support AlwaysOn Availability Groups.
7. Select the Enable AlwaysOn Availability Groups check box, and click OK.
8. Right-click SQL Server (MSSQLSERVER), and select Restart. Close SQL Server Configuration Manager.
9. Repeat steps 1 through 8 for SQL02.

Procedure 22: Add a second virtual network adapter to VMM01 and VMM02

To prevent a single point of failure with the failover cluster, Procedure 22 adds a second NIC to each Virtual Machine Manager cluster node. These steps use 172.16.0.0/24 as the address range; however, you can use any address range that works for your POC.

1. Log on to your MGMT01 host with the contoso\administrator credentials.
2. Open Hyper-V Manager.
3. Register MGMT02 in the console if it's not already registered.
4. Connect to your VMM01 VM and shut it down.
5. Click MGMT02 in the console.

6. Connect to your VMM02 VM and shut it down.
7. Click MGMT01 in the console. Right-click VMM01, and select Settings.
8. In the left pane, click Add Hardware, and select Network Adapter.
9. For a virtual switch, select TenantNetwork vSwitch, and click OK.
10. Repeat steps 7 through 9 for VMM02 on MGMT02.
11. Start both VMM01 and VMM02.
12. Log on to VMM01 using the contoso\administrator credentials.
13. From Server Manager, open the network adapter settings for any adapter. You will see two adapters. Rename the first adapter **TenantNetwork** and rename the second adapter **ClusterNetwork**.
14. Set the following IP address properties for the ClusterNetwork adapter:
 - **IP Address** 172.16.2.1
 - **Subnet Mask** 255.255.255.0
 - **DNS** Leave blank
 - **Default Gateway** Leave blank
15. Save the settings.
16. Log on to VMM02 using the contoso\administrator credentials.
17. From Server Manager, open the network adapter settings for any adapter. You will see two adapters. Rename the first adapter **TenantNetwork** and rename the second adapter **ClusterNetwork**.
18. Set the following IP address properties for the ClusterNetwork adapter:
 - **IP Address** 172.16.2.2
 - **Subnet Mask** 255.255.255.0
 - **DNS** Leave blank
 - **Default Gateway** Leave blank
19. Save the settings.

Procedure 23: Create the guest cluster within VMM01 and VMM02

Virtual Machine Manager is critical to the management of the virtualized infrastructure. From network and storage configuration, through to deployment and management of Hyper-V hosts and VMs, Virtual Machine Manager plays an integral part in the environment. With this in mind, it's imperative to maintain a resilient Virtual Machine Manager configuration. This involves making Virtual Machine Manager highly available at multiple layers—the Virtual Machine Manager management server, the Virtual Machine Manager database, and the Virtual Machine Manager library. The following procedure focuses on installing and configuring the

underlying Windows Server failover cluster, which will support a highly available Virtual Machine Manager management server.

1. Log on to your VMM01 VM with the contoso\administrator credentials.
2. In Server Manager, click Manage, and then click Add Roles And Features.
3. On the Before You Begin Page, click Next.
4. On the Select Installation Type page, select the Role-Based Or Feature-Based Installation option is selected, and click Next.
5. On the Select Destination Server page, select the server location (from a server pool or from a virtual hard disk). After you select the location, select VMM01, and then click Next.
6. On the Select Server Roles page, click Next.
7. On the Select Features page, click Failover Clustering. On the Add Features That Are Required For Failover Clustering page, click Add Features, and then click Next.
8. On the Confirm Installation Selections page, click Install.
9. Repeat steps 1 through 8 on VMM02. When installations are complete on both servers, proceed.
10. On VMM01, open Server Manager, click Tools, and select Failover Cluster Manager.
11. Right-click Failover Cluster Manager, and select Create Cluster.
12. On the Before You Begin page, click Next.
13. On the Select Servers page, add VMM01 and VMM02, and click Next
14. Run all cluster validation tests. A warning that shared storage is not yet configured appears, but the validation tests will pass.
15. For Cluster Name, enter **VMMCluster**, assign IP address **10.10.0.20** (from the TenantNetwork range), and then click Next.
16. A new cluster named VMMCluster.contoso.com should be formed with two nodes: VMM01 and VMM02.

To install Failover Clustering with Windows PowerShell, run the following:

```
#Enable Failover Clustering:  
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools
```

At this stage, the cluster configuration isn't completely redundant because it has no shared storage. A highly available configuration of Virtual Machine Manager doesn't specifically require shared storage for its own configuration. However, when you have a Windows Server failover cluster with an even number of nodes, it's important to have a third vote somewhere in the environment that can be used to avoid issues if one node is down and the other is up.

If you are not familiar with Failover Clustering, the concept of votes (more specifically, the quorum configuration) within a failover cluster determines the number of failures that the cluster can sustain and still remain online. In essence, the quorum configuration affects the

availability of the cluster. A sufficient number of cluster elements must be online or the cluster loses quorum and must stop running. For a deeper understanding of the Failover Clustering quorum configuration, read the following blog post from cluster MVP David Bermingham: http://blogs.msdn.com/b/microsoft_press/archive/2014/04/28/from-the-mvps-understanding-the-windows-server-failover-cluster-quorum-in-windows-server-2012-r2.aspx.

When you deploy this POC's shared storage, you'll be able to adjust your configuration to reflect the additional cluster element. At this point, the system has predetermined the quorum and has no third vote because the environment has no suitable file shares or shared disk available to Virtual Machine Manager to provide this function. This configuration will be revisited when the shared storage is available.

Procedure 24: Create an Active Directory Domain Services container for distributed key management

By default, Virtual Machine Manager encrypts some data in the Virtual Machine Manager database by using the Data Protection Application Programming Interface (DPAPI). For example, Virtual Machine Manager encrypts Run As account credentials and passwords in guest operating system profiles. Virtual Machine Manager also encrypts product key information in virtual hard disk properties for VM role scenarios and configuration.

During the installation of a Virtual Machine Manager management server, you can choose to store the keys to encrypted data on the local computer or configure distributed key management. If you choose to use the local computer, the encryption of this data is tied to the specific computer on which Virtual Machine Manager is installed and the service account that Virtual Machine Manager uses. Therefore, if you move your Virtual Machine Manager installation to another computer, Virtual Machine Manager will not retain the encrypted data. In that case, you must enter this data manually to fix the Virtual Machine Manager objects. Another consideration is with a highly available Virtual Machine Manager management server. If one Virtual Machine Manager server is down, the other Virtual Machine Manager management server, which was previously passive, becomes active, but can't access the local encryption keys that are contained on the inactive server.

In contrast, distributed key management stores the encryption keys in Active Directory Domain Services. Therefore, if you must move your Virtual Machine Manager installation to another computer, Virtual Machine Manager retains the encrypted data because the other computer will have access to the encryption keys in Active Directory Domain Services. This configuration is also a solution for highly available Virtual Machine Manager management servers, and therefore distributed key management is required for these servers.

1. Log on to DC01 with the contoso\administrator credentials.
2. In Server Manager, click Tools, and click ADSI Edit.
3. Right-click ADSI Edit, select Connect To, and select the default naming context.
4. Right-click DC=Contoso, DC=com, and select New Object.

5. Select Container as the object type, enter **CN=VMMDKM,DC=contoso,DC=com** for the container name, click Next, and then click Finish.
6. After you have created the container, the account with which you are installing Virtual Machine Manager must have Full Control permissions to the container in Active Directory Domain Services. Also, the permissions must apply to this object and all descendant objects of the container. If you are installing Virtual Machine Manager as contoso\administrator, you won't need to edit the security properties of the new container.
7. Click OK to close the Container Properties window, and then close ADSI Edit.

Procedure 25: Install Virtual Machine Manager management server prerequisites

The following procedure first configures the necessary prerequisites for the Virtual Machine Manager management servers. This involves installing some of the tools from the Windows Assessment and Deployment Toolkit (ADK), and then some key SQL Server-related utilities, all of which, you downloaded earlier.

1. Log on to your VMM01 VM with the contoso\administrator credentials.
2. Open File Explorer, navigate to \\MGMT01\Software\ADK, and then open Adksetup.
3. On the Specify Location page, keep the default, and click Next.
4. On the Join The Customer Experience Improvement Program page, indicate whether you want to participate, and then click Next.
5. On the License Agreement page, click Accept.
6. On the Select The Features You Want To Install page, check only Deployment Tools and Windows Preinstallation Environment (Windows PE). Keep all other boxes clear. Click Install.
7. When completed, in File Explorer, navigate to \\MGMT01\Software\SQL\, and run the SqlIncli.msi, which installs the SQL Server 2012 SP1 Native Client. Click Next to accept all defaults, click Install, and then click Finish.
8. Run the SqlCmdLnUtils.msi, which installs the SQL Server 2012 SP1 Command Line Utilities. Click Next to accept all defaults, click Install, and then click Finish when the installation is complete.
9. Run Microsoft Update to check for updates. Your WSUS may be configured for downloading updates for only SQL Server 2012 SP1, so check against Microsoft Update to confirm no updates exist for these additional tools. When updating is complete, restart VMM01.
10. Repeat steps 1 through 9 on VMM02.

Procedure 26: Install Virtual Machine Manager management server

The following procedure installs and configures the Virtual Machine Manager server features on both VMM01 and VMM02. The installer will recognize that both of these are part of a failover cluster and customize the installation to reflect this specific deployment option. This ensures you're running on a resilient configuration from the start.

1. On VMM01, navigate to \\MGMT01\Software\VMM\Extracted, and run Setup.exe.
2. On the splash screen, click Install.
3. Click VMM Management Server. The VMM Console check box will automatically be selected. When asked if you want to install in a highly available mode, click Yes.
4. On the Product Registration Information page, enter your relevant information, including the product key, and then click Next. Do not enter a product key if you want to use the product for a 180-day evaluation period.
5. Accept the license agreement, and click Next.
6. On the Customer Experience Improvement Program page, make your selection for participation, and then click Next.
7. On the Microsoft Update page, select On (Recommended), and click Next.
8. On the Installation Location page, select the default location, and then click Next.
9. On the Prerequisites page, the wizard checks for prerequisite hardware and software, and this configuration should pass.
10. On the Database Configuration page, select SQL01, which will be your primary SQL server in your SQL server cluster, accept the other defaults, and then click Next.
11. On the Cluster Configuration screen, enter **VMM-HA** as the name and **10.10.0.21** as the IP address, and then click Next.
12. On the In The Configure Service Account And Distributed Key Management page, enter the domain service account details: **SCVMM_svc** and the appropriate password for the account. Since this will be a highly available configuration, you have no choice but to use Distributed Key Management. In the empty box, enter **CN=VMDKM, DC=contoso, DC=com**, and click Next.
13. On the Port Configuration page, accept the defaults, and click Next.
14. On the Library Configuration page, accept the defaults for now, and click Next. When the resilient file server storage has been configured, you will add further storage to Virtual Machine Manager to hold the important VM-related data.
15. On the Installation Summary page, review the details, and then click Install.

- 16.** When installation is complete, on the Setup Completed Successfully page, clear the Open The VMM Console When This Wizard Closes check box. Click Close.
- 17.** When Windows Update opens, check for updates to Virtual Machine Manager. Install any missing updates, and restart if necessary.

When the first SCVMM cluster node is complete, you can work on the second node. The wizard for the second node will have many items marked as read-only or unavailable because the initial configuration on VMM01 determined these values.

- 1.** On VMM02, navigate to \\MGMT01\Software\VMM\Extracted, and run Setup.exe.
- 2.** On the splash screen, click Install.
- 3.** Select the check box for VMM Management Server. The System Center Virtual Machine Manager console will automatically be selected. The Virtual Machine Manager installer will not only recognize that it is installing on a failover cluster, it will recognize the presence of an existing Virtual Machine Manager configuration on that cluster and will offer to add this server as a node to the existing configuration. Click Yes.
- 4.** On the Product Registration Information page, enter your relevant information, including the product key, and then click Next. Do not enter a product key if you want to use the product for a 180-day evaluation period.
- 5.** Accept the license agreement, and click Next.
- 6.** On the Customer Experience Improvement Program page, make your selection for participation, and then click Next.
- 7.** On the Microsoft Update page, select On (Recommended), and click Next.
- 8.** On the Installation Location page, select the default location, and then click Next.
- 9.** On the Prerequisites page, the wizard checks for prerequisite hardware and software, and this configuration should pass.
- 10.** On the Database Configuration page, the database server is displayed as a read-only value in the Server Name text box. Click Next.
- 11.** On the Configure Service Account And Distributed Key Management page, provide the password of the domain account that the Virtual Machine Manager server will use.
- 12.** On the Port Configuration page, click Next.
- 13.** On the Library Configuration page, click Next.
- 14.** On the Installation Summary page, review the settings, and click Install.
- 15.** When installation is complete, on the Setup Completed Successfully page, clear the Open The VMM Console When This Wizard Closes check box. Click Close.
- 16.** When Windows Update opens, check for updates to Virtual Machine Manager. Install any missing updates, and restart if necessary.

Procedure 27: Create an AlwaysOn Availability Group for VirtualManagerDB

The following procedure sets the recovery model of the Virtual Machine Manager database (VirtualManagerDB) to Full, ensuring the transaction log does not get cleared after every transaction. This is a requirement for AlwaysOn Availability Groups. In a production environment, you will set up a maintenance plan for VirtualManagerDB that protects the database with a differential or incremental backup at least once a day. Perform a full backup at least a few times a week—or whenever your scenario calls for it. You can take this backup from the secondary replica, rather than the primary. After each full backup, you can clear the transaction log to control the backup size. The second thing this procedure does is perform a full backup of VirtualManagerDB. This full backup is also a requirement for performing the initial replication of the database.

1. Log on to your SQL01 VM with the contoso\administrator credentials.
2. Open SQL Server Management Studio. In the Connect To Server window, ensure that SQL01 is listed, and use Windows Authentication to log in. Click Connect.
3. Expand Databases, right-click VirtualManagerDB, and select Properties.
4. Click Options. Change the recovery model to Full, and click OK.
5. Open File Explorer, and navigate to E:\. Create a new folder on E:\AvailGroups.
6. Share E:\AvailGroups with authenticated users as read/write. If you have a specific agent proxy that you use, you can share it directly with the agent proxy account.
7. Right-click VirtualManagerDB, select Tasks, and then select Backup.
8. You can either click Add to create a specific backup device or you can click OK to back up VirtualManagerDB to the default file backup device. Wait for the backup to complete. It should take only a few seconds.
9. In the left pane, expand AlwaysOn High Availability.
10. Right-click Availability Groups, and select New Availability Group Wizard.
11. On the Introduction page, click Next.
12. On the Specify Availability Group Name page, enter **HA_VMM_GROUP**.
13. On the Select Databases page, select the check box next to VirtualManagerDB. Note that this database is marked as meeting the prerequisites. Click Next.
14. On the Specify Replicas page, click Add Replica. In the Connect To Server window, enter **SQL02** next to Server Name, and click Connect.
15. For SQL01 and the newly added SQL02 server instance, select the box for Automatic Failover. This will also automatically select the box for synchronous replication for each SQL Server instance.
16. Click the Listener tab. Select the Create An Availability Group Listener option. For the name, enter **VMMlistener**. For the port, enter **5000**, and ensure Static IP is selected for IP address.

17. To configure the static IP, click Add, select the 10.10.0.0/24 subnet, enter **10.10.0.22** (or whatever free IP address is on the tenant network you prefer), and then click OK.
18. On the Select Initial Data Synchronization page, select Full. The wizard subsequently triggers a full backup of the database, logs to a location you specify, and restores this backup to the secondary SQL Server instance, in this case, SQL02. From there, replication begins synchronously between SQL01 and SQL02.
19. In the Specify A Shared Network Location Accessible By All Replicas text box, enter **\\MGMT01\Software\SQL**, and click Next to begin validation.
20. After validation is complete, review the results, and click Next.
21. On the Summary page, review all of your selections, and click Finish to begin the process. This will take a few moments to complete. When the process is complete, close SQL Server Management Studio.
22. On SQL01, open Failover Cluster Manager, expand SQLCluster.contoso.com, and then click Roles. A new role named HA_VMM_GROUP will be running.

Procedure 28: Finalize Virtual Machine Manager installation

This procedure uses the cluster name to log in to the Virtual Machine Manager management server to check connectivity.

1. On VMM01, log on as contoso\administrator.
2. Double-click the System Center Virtual Machine Manager console icon.
3. For the connection properties, enter **VMM-HA:8100**, which is the highly available Virtual Machine Manager cluster created in Procedure 26.

This page intentionally left blank

Configuring network infrastructure

Logical networks are the first building block in the foundation of network sites, IP address pools, and logical switches, all of which are discussed in this chapter. In this configuration walkthrough, you will configure three logical networks. To these logical networks, Microsoft System Center 2012 R2 Virtual Machine Manager will apply standard configurations and addressing schemas for tenant and datacenter traffic. It is important to define logical networks and IP address schemas for the different types of traffic because it will enable you to easily manage virtual machines (VMs) and host operations at scale.

You start by creating a logical network and IP address pool for the tenant network. You can use a logical network to organize and simplify network assignments for hosts, VMs, and services. In addition, you can use logical networks to describe networks that have different purposes, to create traffic isolation, and even to support traffic that requires different types of service-level agreements (SLAs). As part of logical network creation, you can create network sites to define the Virtual Local Area Networks (VLANs) and IP subnets, and IP subnet and VLAN pairs, and then associate them with the logical network in different physical locations.

A logical network, together with one or more associated network sites, is a user-defined named grouping of IP subnets, VLANs, or IP subnet and VLAN pairs that is used to organize and simplify network assignments.

Configuration walkthrough

In the following sections, you will walk through the key steps required to set up the network infrastructure in System Center Virtual Machine Manager. At this stage, you don't have any hosts under management. But because you're configuring the network infrastructure first, these steps will streamline the process of deploying VMs and service when the time comes.

Figure 3-1 depicts the logical and physical layers of a virtualized networking infrastructure. As you go through this chapter, you'll build each piece of the model and configure the System Center Virtual Machine Manager network fabric that you will use in this proof-of-concept (POC) configuration. This chapter covers the layers in the model in the order that you would

build them, starting with logical networks, rather than working from top to bottom of the diagram in Figure 3-1.

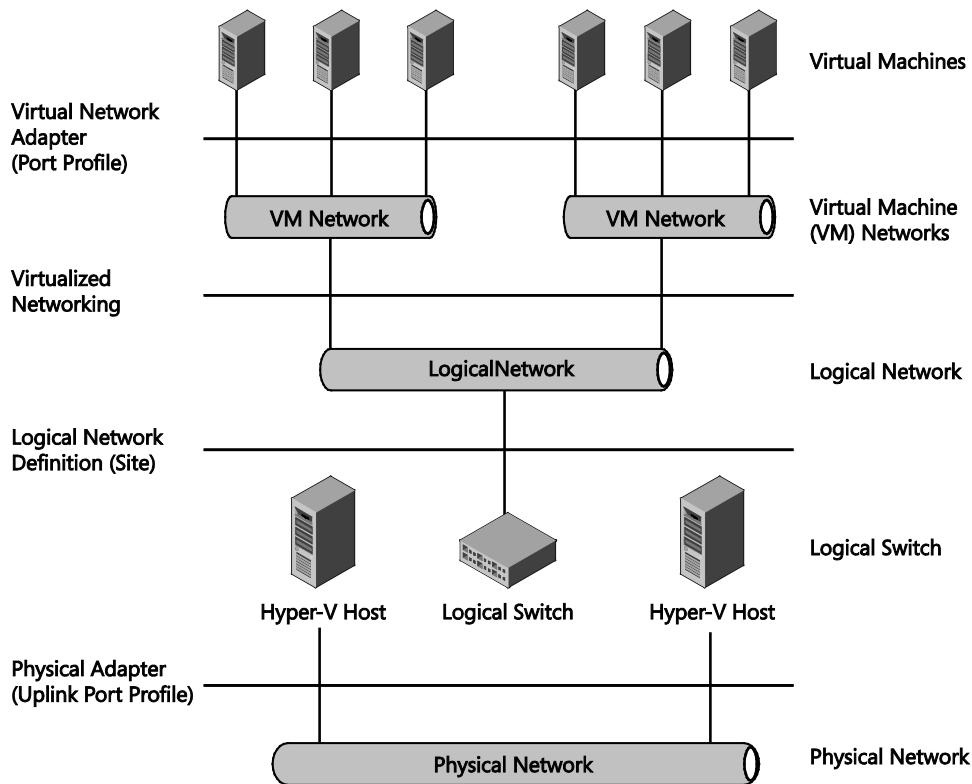


FIGURE 3-1 Logical and physical network layers

Logical networks

A logical network is a user-defined named grouping of IP subnets and VLANs or groupings of pairs of IP subnets and VLANs that is used to identify, organize, and simplify network assignments. Some possible examples include "BACKEND," "FRONTEND," "LAB," "MANAGEMENT," and "BACKUP." Because logical networks represent an abstraction of the underlying physical network infrastructure, they enable you to model the network based on business function and connectivity properties.

After you have created a logical network, you can use it to specify the network on which to deploy a host or a VM (as a standalone or as part of a service). Administrators can assign logical networks as part of VM and service creation without having to understand the network details.

You can use logical networks to designate networks with different purposes, to create traffic isolation, and to provision networks for different types of SLAs. For example, for a tiered application, you can group IP subnets and VLANs that are used for the front-end web tier to define the FRONTEND logical network. You can group IP subnets and VLANs that are used for back-end servers (such as application and database servers) to define the BACKEND logical network. When self-service administrators model the tiered application as a service, by referring to the name associated with the logical network, they can easily pick the logical network that VMs in each tier of the service should connect to.

To deploy VMs and services, you must have at least one logical network. To make a logical network available to a host, you must associate the logical network with a physical network adapter on the host. You create this association for each network adapter.

By default, when you add a Hyper-V host to System Center Virtual Machine Manager management, it automatically creates logical networks that match the first DNS suffix label of the connection-specific DNS suffix on each host network adapter. You'll notice that the first step in this configuration disables this default functionality to ensure that you define all of the logical networks yourself.

Network sites

When you create a logical network, you can create one or more associated network sites. A network site associates one or more subnet and VLAN and pairs of subnets and VLANs with a logical network. You can define the host groups to which the network site is available. For example, if you have a Seattle host group and a New York host group and you want to make the BACKEND logical network available to each, you can create two network sites for the BACKEND logical network. You can scope one network site to the Seattle host group (and any desired child host groups), and you can scope the other network site to the New York host group (and any desired child host groups), adding the appropriate subnets and VLANs for each location, as shown in Figure 3-2.

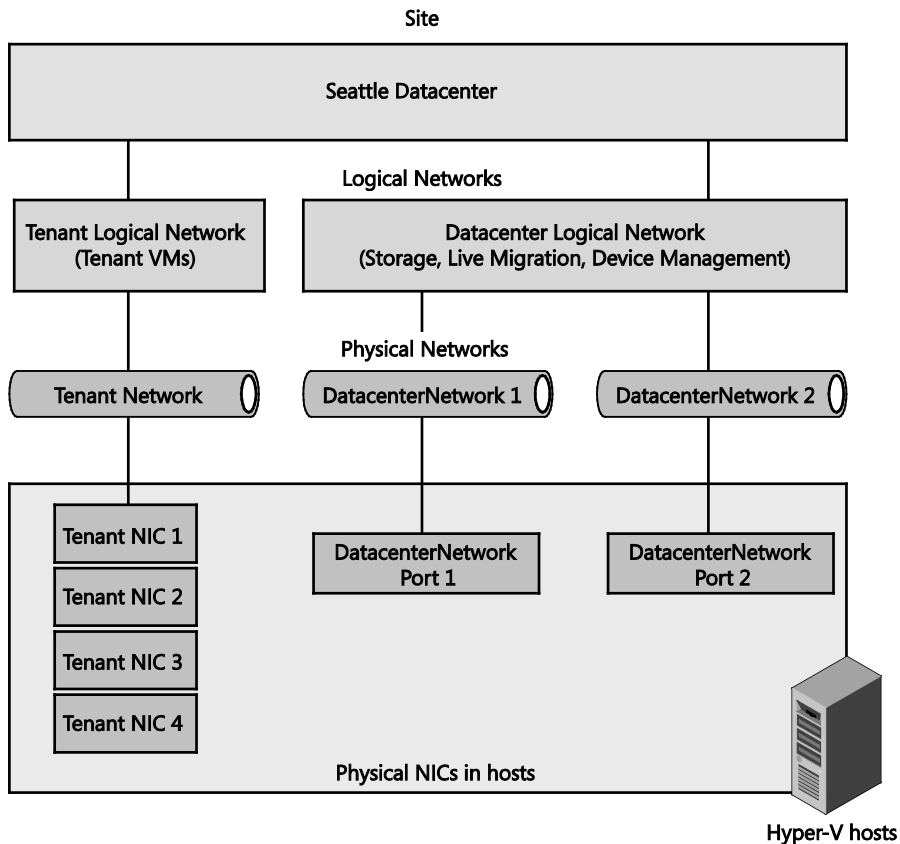


FIGURE 3-2 Logical networks and network sites

Procedure 1: Create a logical network and site for tenant traffic

In this procedure, you create a logical network that will be used for tenant traffic. The logical tenant network will support VM traffic and, later, network virtualization. As part of completing the creation wizard, you'll also create a default network site, which will be associated with the default All Hosts host group.

1. Log on to your VMM01 VM using contoso\administrator credentials.
2. From the desktop, launch the System Center Virtual Machine Manager console, and enter **VMM-HA** as the name. Click Connect. By entering VMM-HA, you'll be connecting to the highly available System Center Virtual Machine Manager configuration you constructed in Chapter 2, "Deploying the management cluster."
3. Open the Settings workspace. Under Settings in the top-left corner, ensure General is selected.

4. In the main console window, double-click Network Settings.
5. In the Network Settings window, clear the Automatic Creation Of Logical Networks check box, shown in Figure 3-3, and then click OK. By clearing this box, you prevent the automatic creation of a logical network for each DNS suffix on a Hyper-V host when the host is added to System Center Virtual Machine Manager's management control. In this POC, you want to keep control over the logical networks that are created, and this setting allows that control.




FIGURE 3-3 The Create Logical Networks Automatically check box

6. Open the Fabric workspace on the lower-left side of the console.
7. At the top of the System Center Virtual Machine Manager console, on the Home tab, in the Show group, click Fabric Resources.
8. In the Fabric pane, expand Networking, and then click Logical Networks.
9. On the Home tab, in the Create group, click Create Logical Network.
10. The Create Logical Network Wizard opens. On the Name page, enter **Tenant_LN**.
11. Leave the One Connected Network option selected, as shown in Figure 3-4, and select the check box next to Allow New VM Networks Created On This Logical Network To Use Network Virtualization. (You will learn more about network virtualization in Chapter 6, "Configuring network virtualization.")
12. Select the check box next to Create A VM Network With The Same Name To Allow Virtual Machines To Access This Logical Network Directly, and then click Next.





FIGURE 3-4 VM network check box settings for a new tenant logical network

13. On the Network Sites page, click Add. System Center Virtual Machine Manager automatically generates a site name that consists of the logical network name followed by an underscore and a number.
14. Under Host Groups That Can Use This Network Site, select All Hosts. For now, you'll assign the logical network and IP address pool to the All Hosts host group because you have yet not created any additional host groups. However, when you start adding compute hosts and management hosts to the System Center Virtual Machine Manager

management scope, you'll assign the logical network and network sites to the particular host groups you create.

15. Under Associated VLANs And IP Subnets, shown in Figure 3-5, click Insert Row. Associate a particular network with the site.
16. Leave the VLAN ID blank, and set the network associated with this site to 10.10.0.0/24 to represent the Tenant Network range. By default, if you leave the VLAN field empty, System Center Virtual Machine Manager assigns a VLAN of 0. This tells System Center Virtual Machine Manager not to use VLANs. In trunk mode, VLAN 0 indicates a native VLAN. You'll define specific IP address pools in a subsequent step. Click Next.



VLAN	IP subnet
	10.10.0.0/24

Network site name: Tenant_LN_0

FIGURE 3-5 Tenant network VLAN and IP subnet settings

17. On the Summary page, review the settings, and then click Finish.
18. The Jobs window appears. Make sure the job has a status of Completed, and then click Close to close the Jobs window.
19. Verify that the logical network appears in the Logical Networks And IP Pools pane. Right-click the logical network, and click Properties. Click the Network Site tab, and verify that the intended network sites appear on the tab. Click Close in the Properties window.
20. Stay on the Networking view within the Fabric workspace of the System Center Virtual Machine Manager console for the next step.

Before you proceed to the next step in the configuration, it's important to understand a little more detail about static IP address pools and how System Center Virtual Machine Manager helps you to manage them.

Static IP address pools

If you associate one or more IP subnets with a network site, you can create static IP address pools from those subnets. Static IP address pools make it possible for System Center Virtual Machine Manager to automatically allocate static IP addresses to VMs that are running on any managed Hyper-V, VMware ESXi, or Citrix XenServer host. From the pool to standalone VMs, System Center Virtual Machine Manager can automatically assign static IP addresses to VMs that are deployed as part of a service and assign IP addresses to physical computers when you use System Center Virtual Machine Manager to deploy them as Hyper-V hosts or scale-out file servers.

Additionally, when you create a static IP address pool, you can define a reserved range of IP addresses that can be assigned to load balancers as virtual IP addresses. System Center Virtual Machine Manager automatically assigns a virtual IP address to a load balancer during the deployment of a load-balanced service tier. You will not be deploying a load balancer as part of this POC.

When you create a static IP address pool, you can configure associated information, such as default gateways, DNS servers, DNS suffixes, and Windows Internet Name Service (WINS) servers. All of these settings are optional. IP address pools support both IPv4 and IPv6 addresses. However, you cannot mix IPv4 and IPv6 addresses in the same IP address pool.

Procedure 2: Create an IP address pool for the Tenant_LN logical network

In this procedure, you'll create a static IP address pool for System Center Virtual Machine Manager to use when assigning IP addresses on this logical network, similar to the way that DHCP assigns addresses. However, on this particular logical network, System Center Virtual Machine Manager uses static IP address pools when assigning IP address information to newly provisioned VMs.

1. While you're still logged in to VMM01, in the Fabric pane, ensure Networking is expanded, and then click Logical Networks.
2. On the Home tab, in the Show group, click Fabric Resources.
3. In the Logical Networks And IP Pools pane, click the Tenant_LN logical network on which you're going to create the IP pool.
4. Right-click the Tenant_LN logical network, and select Create IP Pool. The Create Static IP Address Pool Wizard opens.
5. On the Name page, enter Tenant_LN_Pool and verify that it is being associated with the Tenant_LN logical network. Tenant_LN should automatically be selected in the logical networks drop-down list. Click Next.
6. On the Network Site page, shown in Figure 3-6, ensure that Use An Existing Network Site is selected. The Tenant_LN_0 site should automatically be selected, and it should reflect the network information you provided in the previous step. Click Next.

Specify a network site and the IP subnet

Select an existing network site and IP subnet from the logical network you have chosen or create a new one. Specify the IP subnet using classless inter-domain routing (CIDR) notation; for example 192.168.1.0/24.

Create a multicast IP address pool
 Use an existing network site
 Create a network site

Network site:

IP subnet: VLAN:

Host groups that can use this network site:


 All Hosts

FIGURE 3-6 Network site and IP settings for tenant network

- On the IP Address Range page, under IP Address Range, set the starting IP address to **10.10.0.100**, and set the ending IP address to **10.10.0.250**, as shown in Figure 3-7. Click Next.

IP address range

Specify the range of IP addresses from the subnet to be managed by this pool.

IP subnet: 10.10.0.0/24

Starting IP address:

Ending IP address:

Total addresses: 151

FIGURE 3-7 Tenant network IP address range

Since you have already assigned several addresses from the range for management VMs, clusters, and other addressable resources, start at host address 10.10.0.100 for the IP range. The gateway address is 10.10.0.254, so leave that out of the static IP address pool, as well. Be aware that you can create multiple IP address pools within a subnet. If you create multiple IP address pools within a subnet, the ranges cannot overlap.

- On the Gateway page, click Insert, enter **10.10.0.254** as the Default Gateway that System Center Virtual Machine Manager will assign when distributing the static IP addresses, and then click Next.

9. On the DNS page, shown in Figure 3-8, set the DNS addresses to **10.10.0.11** and **10.10.0.12**. Select 10.10.0.11, click Move Up to make it the first address in the list, and then click Next. Leave the DNS search suffixes blank. You'll set that with Group Policy later.

DNS Server Address	
10.10.0.11	Insert
10.10.0.12	Delete
	Move Up
	Move Down

FIGURE 3-8 DNS server settings for tenant logical network

10. On the WINS page, leave the boxes empty and click Next.
11. On the Summary page, confirm the settings, and then click Finish. The Jobs window appears. Make sure that the job has a status of Completed, and then close the window.
12. Verify that the IP address pool was created: In the Logical Networks And IP Pools pane, expand the logical network where you created the pool. The IP address pool appears under the logical network.

Procedure 3: Create a second logical network and site for the datacenter network

In this procedure, you'll create a second logical network, this time for the physical datacenter network. In Chapter 2, when you configured your management hosts (MGMT01 and MGMT02), the physical datacenter networks were configured to be used on both of the 10-Gbps RDMA-capable network adapters (see Figure 3-9).

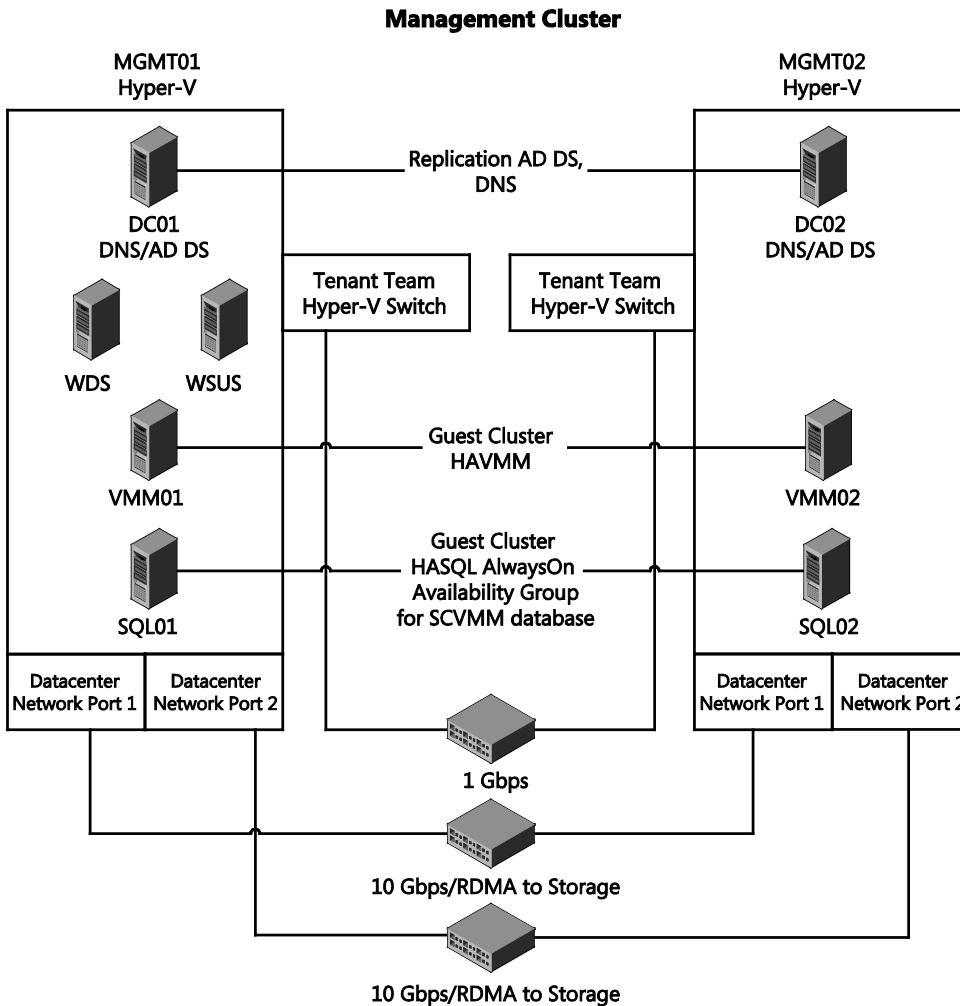


FIGURE 3-9 Logical diagram of TenantNetwork, DatacenterNetwork 1, and DatacenterNetwork 2 logical and physical networks

One of the RDMA-capable network adapters on the management hosts was given an IP address in the 10.10.1.0/24 range, and the other adapter was given an IP address on the 10.10.2.0/24 range. Both of these subnets are within the datacenter network. You will now reflect this configuration through the creation of the datacenter network logical network. The physical servers will use the datacenter network to access shared storage and to enable functionality, such as live migration.

1. In the Fabric/Network workspace of the System Center Virtual Machine Manager console, right-click Logical Networks, and select Create Logical Network.
2. On the Name page, in the Name text box, type **Datacenter_LN**.

3. Leave the One Connected Network option selected, but this time, leave both of the sub-options clear. Because the Datacenter_LN logical network will be used exclusively by the physical compute, storage, and management nodes—and not by the VMs themselves—you do not need to create a corresponding VM network or enable network virtualization. Click Next.
4. On the Network Site page, click Add. Select All Hosts as the host. As you did with the Tenant_LN logical network, you'll assign the logical network and IP address pool to the All Hosts host group because you have not yet created any additional host groups. However, when you start adding compute and management hosts to the System Center Virtual Machine Manager management scope, you'll assign the logical network and network sites to the particular host groups you create.
5. Under Associated VLANs And IP Subnets, click Insert Row, as shown in Figure 3-10. At this point, you'll simply associate a network with the site, just like with the Tenant_LN. The difference in this case is that the Datacenter_LN logical network has two separate networks for redundancy. Therefore, you will add them both.




FIGURE 3-10 VLAN and IP subnet settings for the Datacenter_LN_0 network site

6. Leave the VLAN ID blank, and set the networks associated with this site to **10.10.1.0/24** to represent the DatacenterNetwork 1 range. Add a second row for **10.10.2.0/24** to represent the DatacenterNetwork 2 range, and then click Next.
7. On the Summary page, confirm your settings, and click Finish
8. The Jobs window opens. Monitor the job for successful completion. When the job is completed, close the Jobs window.
9. Return to the Fabric view, and ensure Logical Networks is selected. You should now see Datacenter_LN listed under Logical Networks.
10. Stay in the Network Fabric workspace of System Center Virtual Machine Manager for the next step.

Procedure 4: Create static IP address pools for the Datacenter_LN logical network

In this procedure, following on from the creation of the Datacenter_LN logical network and corresponding network site, you'll create two static IP address pools for System Center Virtual Machine Manager to use when assigning IP addresses on this logical network. As mentioned

earlier, this network will be used for all datacenter traffic, including storage and live migration. In later chapters, you'll assign IP addresses from these pools when provisioning the compute and storage nodes.

1. Right-click the Datacenter_LN logical network, select Create IP Pool, and then click Next.
2. On the Name page, enter **Datacenter_LN_Pool1**. Verify that it is being associated with the Datacenter_LN logical network by confirming that Datacenter_LN has been automatically selected in the logical networks drop-down list. It should be selected in this list because you right-clicked it to start the configuration. Set it if necessary, and then click Next.
3. On the Network Site page, ensure that Use An Existing Network Site is selected, as shown in Figure 3-11. The Datacenter_LN_0 site should automatically be selected, and it should reflect the first network (10.10.1.0/24) you provided in the previous step. If not, select 10.10.1.0/24. Click Next.

Specify a network site and the IP subnet

Select an existing network site and IP subnet from the logical network you have chosen or create a new one. Specify the IP subnet using classless inter-domain routing (CIDR) notation; for example 192.168.1.0/24.

Use an existing network site

Create a network site

Network site: Datacenter_LN_0

IP subnet: 10.10.1.0/24 VLAN: 0

Host groups that can use this network site:

All Hosts

FIGURE 3-11 Network site, IP subnet, and VLAN configuration for the Datacenter_LN_Pool1 IP pool

4. On the IP Address Range page, set the starting IP address to **10.10.1.10** and the ending IP address to **10.10.1.250**, and then click Next.
5. On the Gateway page, click Insert and enter **10.10.1.254** as the gateway, and then click Next.
6. On the DNS page, set the DNS addresses to **10.10.0.11** and **10.10.0.12**. Select 10.10.0.11. Click Move Up to make it the first address in the list, and then click Next.
7. On the WINS page, leave WINS Server empty, and click Next.
8. On the Summary page, review your selections, and then click Finish.
9. The Jobs window opens. Monitor the job for successful completion. When the job is completed, close the Jobs window.
10. Right-click the Datacenter_LN logical network, and select Create IP Pool. You will now create the IP pool for the 10.10.2.0/24 subnet that you defined earlier in the Logical Network Creation Wizard.

11. On the Name page, enter **Datacenter_LN_Pool2**, and verify that it is associated with the Datacenter_LN logical network. Again, confirm that, because you right-clicked on it to start the configuration, Datacenter_LN has been automatically selected in the logical networks drop-down list. Set it if necessary, and then click Next.
12. On the Network Site page, ensure that Use An Existing Network Site is selected. The Datacenter_LN_0 site should automatically be selected. However, you need to use the IP Subnet drop-down list to select subnet 10.10.2.0/24, and then click Next.
13. On the IP Address Range page, set the starting IP address to **10.10.2.10** and the ending IP address to **10.10.2.250**, and then click Next.
14. On the Gateway page, click Insert, enter **10.10.2.254** as the Gateway, and then click Next.
15. On the DNS page, set the DNS addresses to **10.10.0.11** and **10.10.0.12**. Select 10.10.0.11, click Move Up to make it the first address in the list, and then click Next.
16. On the WINS page, leave WINS Server empty and click Next.
17. On the Summary page, review your selections, and then click Finish.
18. The Jobs window opens. Monitor the job for successful completion. When the job is completed, close the Jobs window.

You have now created all of the logical networks that are required for your POC, along with static IP address pools that System Center Virtual Machine Manager will use to supply the compute and storage nodes with IP addresses when you deploy them in subsequent chapters. With these logical networks successfully configured, you are prepared to define the logical switches that will be deployed specifically to your compute nodes in Chapter 5, "Configuring compute infrastructure."

Logical switches

Logical switches allow VMs to communicate out through the physical adapters and NIC teams configured on the Hyper-V host. Logical switches can also enforce standard configurations on Hyper-V host adapters and NIC teams to prevent workloads from experiencing downtime due to misconfigurations.

As Figure 3-12 shows, building a logical switch is like building a layer cake. You start with the Uplink Port Profile, which allows the logical network to communicate with the logical switch. You can configure the logical switch with switch extensions to add functionality, such as packet filtering, forwarding, and third-party management tool integration. You can also configure the logical switch for use with various port classifications. Port classifications allow you to define virtual port profiles for the VMs that use the switch. Settings include the configuration of offload settings such as IPSEC offloading, security settings such as DHCP guard and router guard, and quality of service (QoS) policies.

The logical switch has native uplink port profiles. These profiles add information about the teaming configuration, which logical networks and network sites are available on the physical network adapters, and whether network virtualization is allowed.

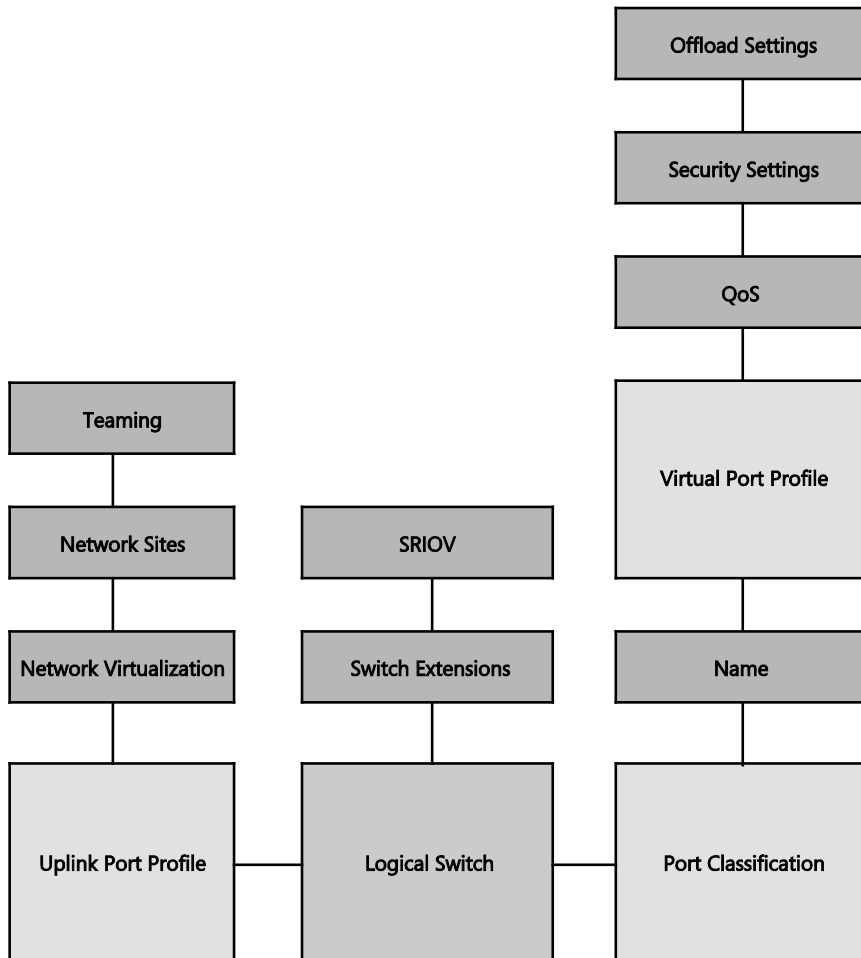


FIGURE 3-12 Building blocks of a logical switch

In Chapter 2, when you deployed and configured your management hosts, MGMT01 and MGMT02, several steps were associated with configuring the network cards on each of the hosts. Specifically, you combined the four 1-Gbps network adapters on each host and used Server Manager to create a NIC team called TenantNetwork Team. As part of this step, you specified a specific teaming mode and load-balancing algorithm. When that was completed, you used Hyper-V Manager on each host to create an external virtual switch (which, is a Hyper-V extensible switch). This external virtual switch was bound to TenantNetwork Team. This binding allows any VMs that are running on the management hosts to communicate out of the host via the external virtual switch, via TenantNetwork Team, onto the physical network.

During that process, all of the work was manual. Performing those actions on tens or hundreds of hosts would be a time-consuming exercise. You could use Windows PowerShell to script the process, but that would require additional time investment to write, test, and then run the scripts. Fortunately, there is a better way. That better way is to use logical switches.

Technically, a logical switch is still the same Hyper-V extensible switch (referred to as the external virtual switch above) that you deployed in the previous chapter. However, a logical switch is a System Center Virtual Machine Manager-only construct, and it wraps the underlying Hyper-V extensible switch with additional granular controls and policy and centralizes the management of the switch within System Center Virtual Machine Manager. This enables you to deploy a logical switch to one or many Hyper-V hosts within the environment from the System Center Virtual Machine Manager console.

Essentially, a logical switch acts as a container for the properties, settings, and capabilities that you want the underlying network adapters to have. In some ways, a logical switch is almost like a network adapter template. You specify the characteristics you'd like as part of the template, and when you deploy the switch, it matches the original template. Instead of configuring individual properties or capabilities for each network adapter, you can specify the capabilities in a logical switch and then apply those capabilities to the appropriate adapters. This can simplify the configuration process dramatically.

Logical switch creation, however, involves a fairly complex wizard, especially the first time you go through the process. One reason it is complex is because a logical switch is made up of building blocks that you must define in advance. Then, to create the logical switch itself, you combine these building blocks to form the logical switch, which you can then deploy to your Hyper-V hosts.

What are these building blocks? Four key building blocks are important to understand, and you will configure them for use in this POC configuration. The four key building blocks, which will be covered in detail in the step-by-step process, are:

- **Hyper-V port profile: Uplink port profile** When you deploy a logical switch to a host (or hosts) you select which physical network adapter (or adapters) you want to bind the logical switch to. When you select the logical switch, System Center Virtual Machine Manager configures the settings for the physical network adapter based on your uplink port profile configuration.
- **Hyper-V port profile: Virtual network adapter port profile** When you deploy a logical switch to a host, the logical switch will enable VMs on that host to communicate out to the physical network. These VMs have one or multiple virtual network adapters (vNICs) that will connect to the logical switch. With the virtual network adapter port profile, you can control the characteristics of that virtual network adapter, controlling the enablement of offload features such as Virtual Machine Queue (VMQ), IPsec Task Offload, and Network QoS, along with several security-specific features. By capturing these settings in a virtual network adapter port profile, in essence, you are creating a template of a vNIC. Upon VM deployment, you can quickly select the virtual network adapter port profile, and all of the relevant

settings for that vNIC will be quickly applied. By adding a virtual network adapter port profile as a building block within the logical switch, you're controlling the network characteristics of the VMs that will attach to that logical switch.

- **Port classification** Port classifications provide global names for identifying different types of virtual network adapter port profiles. You can use a port classification across multiple logical switches, and the settings for the port classification remain specific to each logical switch. For example, you might create one port classification named FAST to identify ports that are configured to have more bandwidth, and another port classification named SLOW to identify ports that are configured to have less bandwidth.
- **Switch extensions** A Hyper-V extensible switch extension is a Network Driver Interface Specification (NDIS) filter or Windows Filtering Platform (WFP) filter that runs inside the Hyper-V extensible switch. Switch extensions, as the name suggests, extend the functionality of the base extensible switch. Two examples of extended functionality are traffic monitoring and virtual firewalls. You don't have to manually deploy these switch extensions to each Hyper-V host. Instead, you can deploy third-party switch extensions at the same time as you deploy the logical switch if you include switch extensions as a building block within the logical switch. However, note that the use of switch extensions is out of scope for this POC configuration guidance.

Procedure 5: Create an uplink port profile

Now that you understand the core building blocks of the logical switch, you can begin to construct those building blocks. In this procedure, you'll create an uplink port profile. As mentioned earlier, when you deploy a logical switch to a Hyper-V host, the uplink port profile determines how the physical network adapter, or adapters, are configured. By capturing these requirements in an uplink port profile, you'll be better able to standardize the configuration of the network across multiple hosts within the environment. In this step, you will be creating a single uplink port profile, but you will also learn more details about scenarios that require additional uplink port profiles.

1. In the Fabric workspace, in the Networking view, in the System Center Virtual Machine Manager console, right-click Port Profiles, and select Create Hyper-V Port Profile.
2. On the General page, in the Name text box, enter **Tenant_LN_UPP**, and select Uplink Port Profile, as shown in Figure 3-13.

Select the type of Hyper-V port profile

You can create virtual network adapter port profiles for use by hosts and virtual machines, or uplink port profiles for use on uplink ports.

Name:

Description:

Type of Hyper-V port profile:

Virtual network adapter port profile

Uplink port profile

Load balancing algorithm:

Teaming mode:

FIGURE 3-13 Load balancing algorithm and teaming mode configuration for the Tenant_LN_UPP Hyper-V port profile

Notice the options available for selection now that you have selected the uplink port profile. You have options for load-balancing algorithm and teaming mode. In the previous chapter, you used Server Manager to create TenantNetwork Team. As part of that wizard, you specified both of these options. By creating an uplink port profile, you can capture the same kind of information. Then, every time this uplink port profile is applied to a network adapter, these settings will also be applied.

This explanation begs the question: What if you have multiple datacenters and a different physical switch is in each datacenter? Perhaps one datacenter has physical switches that support static or Link Aggregation Control Protocol (LACP) teaming, while the switches in the other datacenter need to use the switch-independent teaming mode. To solve this problem, you simply create two uplink port profiles, one for each type of teaming mode required. You assign both uplink port profiles to the logical switch. Then, at logical switch deployment time, you select which uplink port profile to apply to the physical network adapters.

In addition to the teaming mode, you also have the option to specify the load balancing algorithm. As discussed in the previous chapter, you can choose from several algorithms, including Hyper-V port, transport ports, IP addresses, MAC addresses, dynamic, and host default. Leave the selection as Host Default. When this uplink port profile is applied to a host running Windows Server 2012 R2 Hyper-V, the host default is dynamic, thus this is the load balancing algorithm that will be used. If you were applying the same uplink port profile to a Windows Server 2012 Hyper-V host, it would use the Hyper-V port algorithm, since this was the default in that release. Dynamic was introduced in Windows Server 2012 R2.

3. On the General page, click Next.
4. On the Network Configuration page, select the check box next to the Tenant_LN_0 network site. This selection essentially indicates to System Center Virtual Machine Manager that when this particular uplink port profile is attached to a physical network adapter, it can connect to the listed network sites and accompanying logical networks.
5. Leave the Enable Hyper-V Network Virtualization option clear if you're using Windows

Server 2012 R2 because it is always enabled for Windows Server 2012 R2. If you were using Windows Server 2012, you would then select Enable Network Virtualization. You will be using Network Virtualization for this logical network later in the POC configuration. Click Next.

6. On the Summary page, review the selections and configuration, and then click Finish.
7. The Jobs window opens. Monitor the job for successful completion, and then close the Jobs window.

Procedure 6: Create a virtual network adapter port profile

As mentioned earlier, you can use a virtual network adapter port profile to capture a set of specific settings for a VM network adapter and store those settings for future deployment usage. Every time you create a new VM, you'll be able to select a particular virtual network adapter port profile, and System Center Virtual Machine Manager will automatically apply your chosen settings to the vNIC of that VM. This automation will save you time and ensure a standardized approach to vNIC deployment.

In this procedure, you will create a new default virtual network adapter port profile. You'll use it as the standard vNIC configuration for your tenant workloads going forward.

1. In the Fabric workspace, in the Networking view, in the System Center Virtual Machine Manager console, right-click Port Profiles, and select Create Hyper-V Port Profile.
2. On the General page, in the Name text box, enter **Tenant_LN_vNIC**, select Virtual Network Adapter Port Profile, and then click Next.
3. On the Offload Settings page, there are three options. These options relate to features and functionality that require specific hardware capabilities within the underlying physical network adapter. These capabilities are as follows:
 - **Virtual Machine Queue (VMQ)** Packets that are destined for a virtual network adapter are delivered directly to a queue for that adapter, the VMQ. They do not have to be copied from the management operating system to the VM.
 - **IPsec Task Offload** With IPsec Task Offload, some or all of the computational work that IPsec requires for encryption and decryption is shifted from the computer's CPU to a dedicated processor on the network adapter.
 - **Single-Root I/O Virtualization (SR-IOV)** With SR-IOV, a network adapter can be assigned directly to a VM. The use of SR-IOV maximizes network throughput while minimizing network latency and minimizing the CPU overhead that is required to process network traffic. To function, SR-IOV requires support from the host hardware and firmware, the physical network adapter, and drivers in the management operating system and the guest operating system. To function correctly, SR-IOV must be enabled in multiple places—in particular, as part of the virtual network adapter port profile creation process, the logical switch creation process, and finally during logical switch deployment. For this POC configuration, you will not be using SR-IOV.

4. On the Offload Settings page, select Enable Virtual Machine Queue, and then click Next.
5. On the Security Settings page, a number of options can be enabled. Ensure that the Allow Guest Specified IP Addresses option (only available for VMs on Windows Server 2012 R2) is selected because this will be important for later VM deployments. You can leave the other boxes clear; however, for reference, following is an explanation of the remaining options:
 - **Allow Media Access Control (MAC) Spoofing** With MAC spoofing, a VM can change the source MAC address in outgoing packets to an address that is not assigned to that VM. For example, a load-balancer virtual appliance might require this setting to be enabled.
 - **Enable DHCP Guard** With DHCP guard, you can protect against a malicious VM that represents itself as a DHCP server for man-in-the-middle attacks.
 - **Allow Router Guard** With router guard, you can protect against advertisement and redirection messages that are sent by an unauthorized VM that represents itself as a router.
 - **Allow Guest Teaming** With guest teaming, you can team the virtual network adapter with other network adapters that are connected to the same switch.
 - **Allow Electrical and Electronics Engineers (IEEE) Priority Tagging** With IEEE priority tagging, outgoing packets from the virtual network adapter can be tagged with IEEE 802.1p priority. These priority tags can be used by QoS to prioritize traffic. If IEEE priority tagging is not allowed, the priority value in the packet is reset to 0.
6. On the Security Settings page, click Next.
7. On the Bandwidth Settings page, you have three options that are all associated with defining QoS. You can specify the minimum and maximum bandwidth available to the virtual network adapter. The minimum bandwidth, which can be expressed as megabits per second (Mbps) or as a weighted value (from 0 to 100), controls how much bandwidth the virtual network adapter can use in relation to other virtual network adapters.

For this POC configuration, you can leave the minimum as 0 Mbps. For maximum, enter **1024** Mbps. This will ensure that very noisy VMs don't consume all of the bandwidth available across the tenant network. For a POC, this value is fine; however you may choose to adjust this figure for your specific environment. The setting that's right for you will take into consideration how many VMs you allow to share the channel and what measure of bandwidth must be made available to each VM to ensure a consistent QoS is met. Click Next.
8. On the Summary page, review all the settings and selections you made, and then click Finish.
9. The Jobs window opens. Monitor the job through its completion, and then close the Jobs window.

Procedure 7: Create a port classification

A port classification provides a global name for identifying different types of virtual network adapter port profiles. As a result, a classification can be used across multiple logical switches, while the settings for the classification remain specific to each logical switch. For example, you might create one port classification named FAST to identify ports that are configured to have more bandwidth and one port classification named SLOW to identify ports that are configured to have less bandwidth. You can use the port classifications that are provided in System Center Virtual Machine Manager, or you can create your own port classifications.

In this procedure, you will create a new port classification to represent the previously created virtual network adapter port profile. As building blocks, the port classification and virtual network adapter port profile are technically unrelated. However, especially in self-service environments, when you combine them as part of the logical switch creation process, the port classification provides a meaningful way to associate a name with underlying virtual network adapter characteristics and capabilities.

1. In the Fabric workspace, in the Networking view, in the System Center Virtual Machine Manager console, right-click Port Classifications, and select Create Port Classification.
2. In the Create Port Classifications Wizard, in the Name text box, enter **Default Tenant vNIC**, and then click OK.
3. Observe that the new port classification appears in the list with several existing port classifications. As mentioned earlier, unless you include the port classifications within a logical switch, they don't provide any functionality as a standalone building block.

Procedure 8: Create the logical switch

At this stage, all of the fundamental building blocks required for the creation of the logical switch are in place. As mentioned earlier, the logical switch brings port profiles, port classifications, and switch extensions together so that you can apply them consistently to network adapters on multiple host systems.

1. Within the System Center Virtual Machine Manager console, in the Fabric workspace, in the Networking view, right-click Logical Switches, and select Create Logical Switch.
2. On the Getting Started page, read the information provided. This information outlines what you have covered so far. It provides four key tasks that you must perform before you create the logical switch. These four tasks map nicely to the building blocks outlined earlier: the logical networks and sites; switch extensions (out of scope for this POC configuration); uplink port profiles and virtual adapter port profiles; and (although not called out in the text of this POC) port classifications. After you've read the information, click Next.
3. On the General page, in the Name text box, enter **Tenant_LN_LS**. Note the check box for Enable Single-root I/O virtualization (SR-IOV). Building upon the information discussed earlier, if you were choosing to take advantage of SR-IOV capabilities, this is

another point at which you would enable SR-IOV. During deployment of a logical switch to a Hyper-V host, System Center Virtual Machine Manager enables the specific SR-IOV functionality defined in the virtual port profile during the creation of that logical switch. This is the only time SR-IOV functionality can be enabled. For this POC configuration, leave the selection empty. Click Next.

4. On the Extensions page, any Hyper-V extensible switch extensions that had already been imported into System Center Virtual Machine Manager management appear. Switch extensions (which you can install on the System Center Virtual Machine Manager management server and then include in a logical switch) allow you to monitor network traffic, use QoS to control how network bandwidth is used, enhance the level of security, or otherwise expand the capabilities of a switch. System Center Virtual Machine Manager supports four types of switch extensions:
 - Monitoring extensions can monitor and report on network traffic, but they cannot modify packets.
 - Capturing extensions can inspect and sample traffic, but they cannot modify packets.
 - Filtering extensions can block, modify, or defragment packets. They can also block ports.
 - Forwarding extensions can direct traffic by defining destinations, and they can capture and filter traffic. To avoid conflicts, only one forwarding extension can be active on a logical switch.

For this POC, you have only the default switch extensions, one of which will be selected already. Verify that Microsoft Windows Filtering Platform is selected. When the logical switch is deployed to a host, this particular extension will be deployed with it, saving you administrative effort it would take to enable the capability on each switch on each Hyper-V host.

5. On the Uplink page, shown in Figure 3-14, from the Uplink Mode drop-down list, select Team. Click Add, and under Uplink Port Profiles, select Tenant_LN_UPP. Click OK.

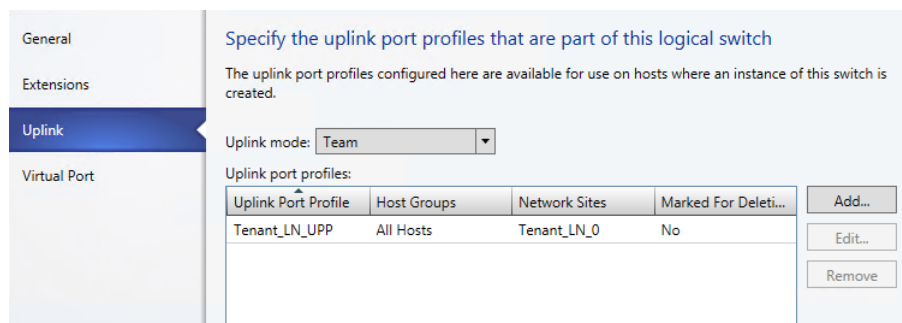


FIGURE 3-14 Specify the uplink port profile used for the logical switch

This is an incredibly important part of the process and one that is pivotal for comprehension. By using the drop-down list to select Team, what you're essentially indicating to System Center Virtual Machine Manager is that when this logical switch is deployed, it will be deployed only onto hosts where multiple physical network adapters will be used for the underlying traffic.

Think about the management hosts, MGMT01 and MGMT02. You used Server Manager to create TenantNetwork Team manually, binding that team to four 1-Gbps physical network adapters as part of the process. You chose a teaming approach for the aggregation of bandwidth that you get by combining four 1-Gbps adapters. But in addition, you got the extra levels of resilience and redundancy of an active/active four-way network adapter team.

By selecting Team from this drop-down list, you're essentially indicating that when you deploy this logical switch to allow VMs to communicate out onto the physical network, you want to ensure they are connected in a redundant manner. Therefore, because you selected Team from the drop-down list and saved this particular logical switch configuration, when you deploy the logical switch in a later chapter, the deployment wizard will specifically ask for multiple network adapters.

System Center Virtual Machine Manager will then team those physical adapters automatically. But how will it know which teaming mode to use or which load balancing algorithm? Recall that you defined that information in the uplink port profile earlier. This is why you added it. At deployment time, System Center Virtual Machine Manager will require multiple network adapters and will ask you for a specific uplink port profile to apply to the team it's about to create. You'll see that behavior later when you deploy the compute nodes.

6. On the Uplink page, click Next.
7. On the Virtual Port page, specify the port classifications and virtual ports that will be used by VMs that will be attached to this logical switch. To do so, click Add. The Add Virtual Port window opens. In this window, you can connect two of the key building blocks: the port classification (which describes what the virtual port can do) and the virtual network adapter port profile (which contains the actual functionality that will be applied to the virtual network adapters). You constructed both of these pieces earlier. First, next to Port Classification, click Browse.
8. In the Select A Port Profile Classification window, select Default Tenant vNIC, and click OK.
9. Return to the Add Virtual Port window, and at the bottom, select the Include A Virtual Network Adapter Port Profile In This Virtual Port check box.
10. From the drop-down list next to Native Virtual Network Adapter Port Profile, select the virtual network adapter port profile you created earlier, called Tenant_LN_vNIC, and then click OK. Because you combined these two objects, when administrators are deploying VMs and they select the port classification for their chosen VM vNIC, System

Center Virtual Machine Manager will automatically apply the characteristics and features contained within the Tenant_LN_vNIC virtual network adapter port profile. (Recall that this will enable VMQ, allow guest-specified IP addresses, and restrict the maximum bandwidth to 1024 Mbps.)

- 11.** Return to the Virtual Port page, select your Default Tenant vNIC from the main window, and on the right side, click Set Default. Click Next.
- 12.** On the Summary page, review the settings, and then click Finish.
- 13.** The Jobs window appears. Monitor the job through successful completion, and then close the Jobs window. The logical switch creation process is now complete.
- 14.** Return to Fabric view, and under Networking, under Logical Switches, your newly created Tenant_LN_LS logical switch will be listed.

This page intentionally left blank

Configuring storage infrastructure

For this proof-of-concept (POC) configuration, you ultimately construct a software-defined storage solution that is underpinned by low-cost, high-volume hardware and transformed through software. Specifically, you take advantage of software features and functionality that come with Windows Server and are enhanced through System Center. At this point, you have defined the physical and logical networking infrastructure in System Center Virtual Machine Manager. The next step is to use the powerful capabilities within System Center Virtual Machine Manager to deploy and configure the storage infrastructure.

You begin by physically connecting your storage hardware. This configuration uses two x86 physical servers, which will become two file server nodes in your Scale-Out File Server (SOFS) configuration. Connected to these two physical servers, by means of industry-standard serial attached SCSI (SAS), will be two JBOD (just a bunch of disks) enclosures. Not only will these two JBODs be connected to the two SOFS nodes, but they will also be connected to one another to aggregate the disks and provide an extra level of redundancy.

The JBODs are in a Storage Area Network (SAN)-like chassis that contains a mixture of hard disk drive (HDD) and solid state drive (SSD) disks, without the intelligence wrapper that typically comes with a SAN. When presented to Windows Server, a JBOD with 24 drives, for example, typically appears in Disk Management as a listing of 24 individual drives that are simply passed through as raw disks from the JBODs to Windows Server. At that point, it's up to Windows Server to transform these into enterprise-class storage. You'll learn more about this process as you go through this chapter.

When everything in your rack is connected, you'll use System Center Virtual Machine Manager to deploy an operating system image to both of the SOFS nodes. From there, System Center Virtual Machine Manager will automate the construction of the SOFS from these freshly deployed hosts. You'll then use the core tools within System Center Virtual Machine Manager and Windows Server Server Manager to aggregate your raw disks into clustered storage pools and slice them into resilient virtual disks with accompanying continuously available file shares. At that point, you'll be ready to present them to the compute nodes, which you'll be deploying in the next chapter.

Scale-Out File Server

Before you move ahead with the deployment, it's a good idea to understand what a SOFS is and what the architecture of the final SOFS will look like when complete. Scale-Out File Server is a Windows Server feature that is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provides the ability to share the same folder from multiple nodes of the same cluster. In this POC configuration, the file-based server application storage includes virtual machines (VMs). With SOFS, you can access a given file share from any of the nodes in the cluster. This ability means that SOFS is perfect as a storage back end for Hyper-V VMs.

You might be wondering whether the term "Scale-Out File Server" means the same as "clustered file server" and whether you can use these terms interchangeably. The answer is no. The clustered file server functionality has been supported in Windows Server since the introduction of Failover Clustering. This type of clustered file server, and therefore all the shares associated with the clustered file server, is online on one node at a time. This is sometimes referred to as active-passive or dual-active. File shares associated with this type of clustered file server are called clustered file shares. This is the recommended file server type when deploying information worker scenarios.

Windows Server 2012 introduced SOFS. This feature lets you store server application data, such as Hyper-V VM files, on file shares and obtain a similar level of reliability, availability, manageability, and high performance that you would expect from a SAN. All file shares are simultaneously online on all nodes. File shares associated with this type of clustered file server are called scale-out file shares. This is sometimes referred to as active-active. This is the recommended file server type when deploying either Hyper-V over Server Message Block (SMB) or Microsoft SQL Server over SMB.

Figure 4-1 shows how everything connects together to form the solution, including two x86 servers, which will be the SOFS nodes. Although this configuration uses two nodes, an SOFS configuration can support up to eight nodes, providing even greater levels of redundancy. However, eight nodes also require additional bandwidth and throughput to the backend disks.

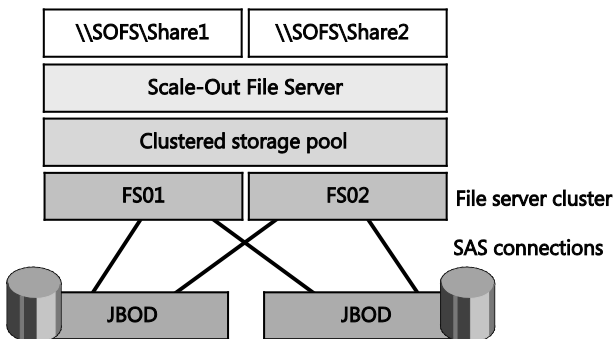


FIGURE 4-1 Logical diagram of a Scale-Out File Server

The back-end disks are enclosed within two 12-bay JBODs that are connected to each other and the SOFS nodes by means of industry-standard SAS cables. Both SOFS nodes will see all 24 disks across the two JBODs. These disks will be a mixture of HDD and SSD so that you can take advantage of the integrated storage tiering and achieve higher levels of performance.

With the SOFS constructed, consumers of the SOFS—specifically, the compute nodes and the existing management cluster, which you will soon construct for this POC—will be able to connect over the 10-Gbps datacenter network. This connection will use the SMB protocol to deliver optimal performance. You'll learn more about this as you progress through this chapter.

Configuration walkthrough

In the following sections, you walk through the key steps outlined in the previous section, building on the networking setup you completed in the previous chapter. This enables you to streamline and centralize the deployment of the SOFS nodes from System Center Virtual Machine Manager.

NOTE The SOFS nodes use two identically configured Dell PowerEdge R620 servers with the following specifications:

- 64-GB RAM, dual Intel Xeon E5-2650 @ 2 Ghz, each with eight cores
- Two dual-port 6-Gbps SAS host bus adapters (HBAs)
- Two 10-Gbps RoCE-capable NICs, used for DatacenterNetwork 1 and DatacenterNetwork 2
- One baseboard management controller (BMC) port for lights-out management over the network

The specifications of your servers might be different from those of this POC, but as a minimum you need two modern CPUs with at least four cores. The SOFS does not need the most powerful CPUs because most traffic is handled by Remote Device Memory Access (RDMA) over Converged Ethernet (RoCE)-capable network cards that process network traffic directly. For a POC configuration, your SOFS nodes do not require a large amount of physical memory, because the SOFS uses storage tiering, which prevents the usage of a feature known as Cluster Shared Volumes (CSV) Cache. CSV Cache is typically one of the largest consumers of RAM on a SOFS node. In this configuration, our servers have 64-GB RAM. From a local storage perspective, your SOFS nodes need at least two HDDs or SSDs that are set up in a RAID-1 configuration and that use the onboard RAID controller. This mirror holds the operating system installation on each node.

Within each SOFS node, you need a minimum of two identical dual-port 6-Gbps SAS HBAs. The specific number you need ultimately depends on how many JBODs you intend to connect to your SOFS, as well as the number of SOFS nodes you have. However, for this configuration, with two SOFS nodes and two JBODs to ensure redundancy and to provide multipathing support, the requirement is two dual-port SAS HBAs. The configuration, as shown in Figure 4-1, ensures you maximize throughput and provides redundant paths.

NOTE The POC JBODs are two identical Dell PowerVault MD1200s with the following specifications:

- Four 250-GB SSD drives for performance per JBOD
- Eight 1-TB 7,200-RPM HDD drives for capacity per JBOD

The specifications of your JBODs might differ from the POC specifications, but as a minimum, your JBOD enclosures should be certified for use with Windows Server Storage Spaces, as indicated on the Windows Server Catalog.

The JBODs should be identically configured. When determining the number of SSDs and HDDs you require, the general rule is to use a 1:4 ratio. The number of JBODs you require is driven by your capacity, performance, and redundancy needs. With two JBODs, the POC configuration can handle a maximum of two disks failing (total, across both JBODs, depending on disk mirroring/parity) before the configuration becomes unavailable. If you add more JBODs, the configuration can handle a greater number of disk failures or complete enclosure failures. You will learn more about enclosure awareness later in this chapter.

A typical production solution assumes a larger number of enclosures with more disks than in this POC. For example, a configuration with 48 7,200-RPM HDDs per JBOD means a total of 192 HDDs across four JBOD enclosures. By choosing 7,200-RPM HDDs, you benefit from huge capacity while consuming less power and at lower cost than higher rotational speed HDDs. In this solution, 7,200-RPM drives provide good performance when matched with a sufficient number of SSDs.

If you have 48 HDDs, using the ratio of 1:4 mentioned earlier, 12 SSDs is an appropriate choice. The storage configuration uses SSDs to create a faster storage tier for frequently accessed data. It also uses SSDs for a persistent write-back cache that reduces the latency of random writes.

When you use 4-TB HDDs and 200-GB SSDs in four 60-bay JBODs, this solution provides approximately 724 TB of raw storage pool capacity per SOFS. After you factor in resiliency and free space for rebuilding storage spaces, this yields roughly 226 TB of space for compute and management VMs.

It's important to note that all disks must be dual-port SAS disks. This ensures that each disk has a connection to all nodes of the SOFS through SAS expanders included in the JBODs.

Procedure 1: Rack and connect SOFS nodes

In this procedure, you physically deploy the hardware into the environment. This might involve the acquisition of new hardware or repurposing of existing hardware. If you already have hardware racked and cabled in your environment, ensure that the networking is configured as specified in this procedure, following the instructions below. The two SOFS nodes are named FS01 and FS02.

1. Rack both physical servers that will become FS01 and FS02.
2. Connect power. At this point, you are just plugging in the hardware. The network

configuration will be completed automatically at deployment time. Connect the network, as follows:

The SOFS nodes will use two separate networks:

- DatacenterNetwork Port 1 (10 Gbps, not teamed)
- DatacenterNetwork Port 2 (10 Gbps, not teamed)

Plug all NICs on each node into physical network ports across separate switches:

- DatacenterNetwork Port 1 on each node to a switch that is RoCE capable
- DatacenterNetwork Port 2 on each node to another switch that is also RoCE capable

Procedure 2: Configure Baseboard Management Controllers

In Chapter 2, “Deploying the management cluster,” you successfully configured System Center Virtual Machine Manager running on top of the management cluster. When you started Chapter 2, System Center Virtual Machine Manager was not available. As a result, the steps for deploying the management nodes was manual. Now that System Center Virtual Machine Manager is running, you can use some of its built-in functionality to streamline deployment of physical servers.

In this chapter, as mentioned earlier, you use System Center Virtual Machine Manager to deploy the SOFS. A key enabler of this bare-metal deployment capability is the use of the Baseboard Management Controller (BMC). The BMC enables out-of-band management. Specifically, the BMC allows System Center Virtual Machine Manager to power on and discover the physical server, control the network boot, and ultimately power off the server, all before the installed operating system starts. Table 4-1 shows the BMC settings you’ll need for FS01 and FS02.

TABLE 4-1 BMC settings for FS01 and FS02

PHYSICAL NIC	PURPOSE	NAME
FS01 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.1 SM: 255.255.255.0 DG: 10.10.0.254
FS02 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.2 SM: 255.255.255.0 DG: 10.10.0.254

NOTE Each hardware vendor’s BMC configuration utility might differ from that of other vendors, so refer to your chosen hardware vendor’s guidance on how to configure your BMC. After it’s configured, make a note of the username, password, and IP address.

For a successful bare-metal deployment, your SOFS nodes must support one of the following out-of-band management protocols:

- Intelligent Platform Management Interface (IPMI) version 1.5 or 2.0
- Data Center Management Interface (DCMI) version 1.0
- System Management Architecture for Server Hardware (SMASH) version 1.0 over WS-Management (WS-Man)

Whichever protocol you use, ensure you use the latest version of firmware for the BMC model.

Procedure 3: Rack and connect JBODs to SOFS nodes

In this procedure, you physically deploy the JBODs and use SAS cables to connect them to the SOFS nodes. This might involve the acquisition of new hardware or repurposing of existing hardware. If you already have hardware racked and cabled in your environment, ensure that the SAS cabling matches the diagram in Figure 4-2. Multipathing solutions use redundant physical path components—adapters, cables, and switches—to create logical paths between the server and the storage device. In the event that one or more of these hardware devices fails, causing the path to fail, multipathing logic uses an alternate path for I/O so that applications can still access their data. To support this configuration in Windows Server, Multipath I/O (MPIO) is enabled as part of the deployment.

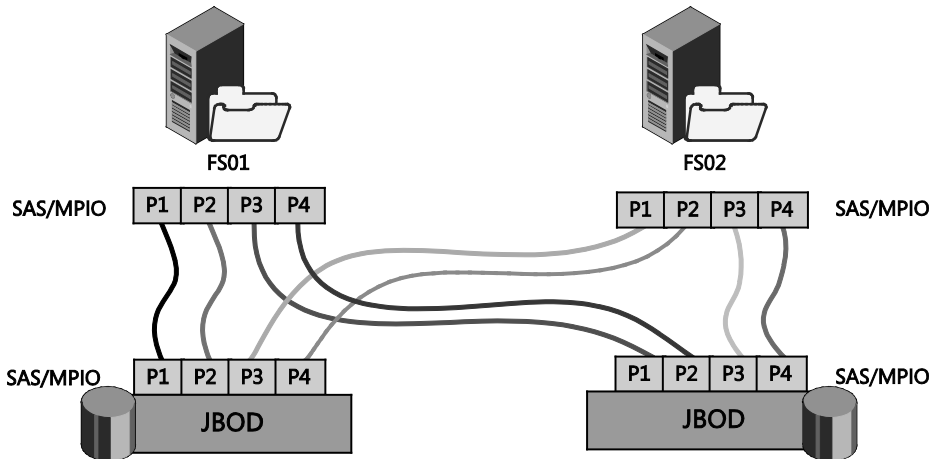


FIGURE 4-2 Two servers, each with two SAS HBAs in failover cluster configuration with two JBODs and MPIO

As shown, this configuration ensures that each of the two SOFS nodes has a connection to each of the controller modules across both JBODs. When connectivity is completed, the JBODs can be powered on.

Procedure 4: Configure a BMC administrator in System Center Virtual Machine Manager

For System Center Virtual Machine Manager to successfully control the SOFS nodes by means of BMCs, you must provide the specific credentials for the BMC for the purpose of bare-metal deployment. You can do this in System Center Virtual Machine Manager by specifying a new run-as account that contains the credentials required for the BMC.

1. Log on to your VMM01 VM using contoso\administrator credentials.
2. From the desktop, launch the System Center Virtual Machine Manager console. For the name, type **VMM-HA**, and click Connect. By entering VMM-HA, you'll be connecting to the highly available System Center Virtual Machine Manager configuration you constructed in Chapter 2.
3. In the bottom-left corner of the System Center Virtual Machine Manager console, click Settings.
4. Expand Security, and select Run As Accounts. On the top ribbon navigation, click Create Run As Account.
5. In the Create Run As Account Wizard, in the Name text box, type **BMC Administrator**. Optionally, you can add a description.
6. In the Username text box, type the user name that is relevant for your BMC settings.
7. Enter a password and password confirmation in the corresponding boxes.
8. Unless your BMC account is within your domain, clear the Validate Domain Credentials check box, and click OK.

You now have a run-as account that corresponds to the BMC on your SOFS nodes. If your SOFS nodes have BMC credentials that differ from one another, you will need to repeat steps 1 through 8 in Procedure 4 to add another run-as account.

Procedure 5: Create a run-as account for SetupAdmin

With your BMC run-as account configured, it's important to ensure that you have another account with enough privilege to perform other administrative actions. A best practice is to create a specific account in Active Directory for this purpose, but for this configuration, use the contoso\administrator account.

1. In the System Center Virtual Machine Manager console, in the bottom-left corner, click Settings.
2. On the Home tab, click Create Run-as Account.
3. Name the run-as account **SetupAdmin**, add a description, and use the contoso\administrator credentials. Click OK, as shown in Figure 4-3.

The image shows a configuration dialog box for a run-as account. It has the following fields and options:

- Name:** SetupAdmin
- Description:** SCVMM Setup Run-as Account
- User name:** Contoso\Administrator
Example: contoso\domainuser or localuser
- Password:** [masked with dots]
- Confirm password:** [masked with dots]
- Validate domain credentials

FIGURE 4-3 SetupAdmin run-as account

Procedure 6: Obtain a virtual hard disk for server deployment

As part of the bare-metal deployment process, rather than installing an operating system to the physical server in the traditional way, System Center Virtual Machine Manager instead deploys a generalized virtual hard disk (VHDX) to the target machines and configures the target host to natively boot from the VHDX. In the following procedure, you obtain the VHDX that is ready for bare-metal deployment.

There are multiple ways to construct a VHDX that can be used for the physical server deployment. One way is to create a new Generation 1 VM, install an appropriate operating system, and then use sysprep with the /generalize and the /oobe (out of box experience) options configured. The resultant VHDX file is then suitable for use as part of a bare-metal deployment.

Another option is to use the Windows Server 2012 R2 ISO file that you downloaded in Chapter 2, along with the Convert-WindowsImage.ps1 Windows PowerShell-based tool, which you can find on the TechNet galleries at <https://gallery.technet.microsoft.com/scriptcenter/Convert-WindowsImageps1-0fe23a8f>. This option quickly creates a VHDX file from the Windows Image (WIM) on the ISO.

However, in Chapter 2, you created a Sysprep.vhdx file that is perfectly suited to this deployment. Use this VHDX file as the basis for physical server deployment. But first, you need to move the VHDX file into your System Center Virtual Machine Manager library.

The challenge at this point, however, is that your System Center Virtual Machine Manager configuration is running as a highly available configuration. Because of this, as part of the System Center Virtual Machine Manager setup wizard, a default library was not created. As a temporary solution, you will use MGMT01 as a library location for the bare-metal deployment phase. In Chapter 5, "Configuring compute infrastructure," when you have your SOFS fully deployed and configured, you will remove MGMT01 and replace it with a dedicated, highly available library server.

1. In the System Center Virtual Machine Manager console, in the bottom-left corner, click Fabric.

2. Expand Infrastructure, and then expand Library Servers.
3. Right-click Library Servers, and select Add Library Server, as shown in Figure 4-4.




FIGURE 4-4 Add Library Server context menu

4. On the Enter Credentials page, click Browse, and select the SetupAdmin run-as account. Click OK, and then click Next.
5. On the Select Library Servers page, type **MGMT01**, and click Add. MGMT01 appears in the list of selected servers. Click Next.
6. On the Add Library Shares page, select the Exports share, and click Next.
7. On the Summary page, click Add Library Servers.
8. Monitor the job for successful completion.
9. When the job is completed, your newly added Library Server appears in the main console window with a status of Responding.
10. In the bottom-left corner of the console, click Library.
11. Expand Library Servers, expand mgmt01.contoso.com, expand Exports, expand Sysprep, and, finally, expand Virtual Hard Disks.
12. In the center pane, right-click SYSPREP.vhdx, and click Properties.
13. In the Properties window, from the Operating System drop-down list, select Windows Server 2012 R2 Standard or Datacenter, depending on which you chose for creation in Chapter 2. Click OK.

NOTE In Chapter 2, when you originally created this Sysprep.vhdx file, the maximum file size for this dynamic VHDX file was set to 80 GB. As part of the bare-metal deployment process, you will have the option to expand this VHDX file from its current size as a dynamic VHDX to its fully expanded size as a fixed VHDX. If your target SOFS nodes do not have enough local storage to store this fixed VHDX size plus the current size of the dynamic VHDX file, plus the size of the system page file, use an alternative, smaller VHDX file. Alternatively, you can choose not to expand the VHDX from dynamic to fixed as part of the deployment process.


Procedure 7: Create a physical computer profile

As of System Center 2012 R2 Virtual Machine Manager, physical computer profiles replace host profiles. The concept of using a physical computer profile for physical server deployment is similar to the concept of using VM templates for VM deployment. Physical computer profiles

define the standardized characteristics of a physical server deployment and let you provision a computer into either a Hyper-V host or a SOFS cluster. Physical computer profiles include configuration settings such as the location of the operating system image to use during host deployment, together with configuration settings for the hardware and operating system. This is a requirement for bare-metal provisioning.

The following procedure describes how to create a physical computer profile in the System Center Virtual Machine Manager library. You will use this profile to provision your SOFS.

1. In the bottom-left corner of the System Center Virtual Machine Manager console, click Library.
2. Right-click Physical Computer Profile, and select Create Physical Computer Profile.
3. On the Profile Description page, shown in Figure 4-5, name the profile **File Server**, add a description, select Windows File Server, and then click Next.



Provide a name for the physical computer profile


Name: File Server

Description: Physical Computer Profile for File Server Cluster Node

Role: VM Host
 Windows File Server

FIGURE 4-5 Defining the physical computer profile

4. On the OS Image page, shown in Figure 4-6, click Browse, and select the VHDX file that you published to the System Center Virtual Machine Manager library in the previous procedure. If desired, select the Do Not Convert The VHD Type To Fixed Type During Deployment check box, and then click Next.



Server role: Windows File Server

VHD file: \\mgmt01.contoso.com\Exports\SYSPREP\Virtual Hard Disks\SYSPREP.vhdx

Virtual hard disk type: Dynamic

Expanded size: 127.00 GB

Current size: 8.63 GB

Minimum partition size needed: 135.63 GB

Do not convert the VHD to fixed type during deployment

FIGURE 4-6 The VHD file to be deployed as part of the WDS deployment

5. On the Hardware Configuration page, under Physical NIC #1, click IP Configuration.
6. Select Allocate A Static IP Address From The Following Logical Network check box, and select DataCenter_LN from the Logical Network drop-down list, as shown in Figure 4-7.




FIGURE 4-7 IP configuration of a physical computer profile

7. At the top of the window, click Add, and then select Physical Network Adapter.
8. Repeat steps 5 through 6 of this procedure for Physical NIC #2.
9. Under Disk And Partitions, click OS, and review the settings. You will be using 100 percent of the available local disk for the deployment, but you can adjust if required.
10. Under Driver Options, click Driver Filter. System Center Virtual Machine Manager allows you to include hardware-specific drivers as part of the deployment process. This section of the wizard assumes you have already imported the drivers into the System Center Virtual Machine Manager library. This will not be covered as part of this configuration. You will use the drivers that are included with Windows Server for the time being. Click Next.
11. On the OS Configuration page, type **contoso.com** as the domain to join, and select the SetupAdmin run-as account as the credentials used to join the domain.
12. For the Admin Password, type the standard password you're using for the POC. This will be used to set the local administrator account password on the server.
13. For Identity Information, enter your company information (optional).
14. For Product Key, enter your product key if you are using one for the POC (optional).
15. Select your time zone.
16. Select an answer file if you're using one for any of your Windows settings (optional).
17. Click Next, and then click Finish.
18. Monitor the job for successful completion, and then close the Jobs window.

Procedure 8: Add a PXE server to System Center Virtual Machine Manager

For System Center Virtual Machine Manager to deploy an operating system to a physical bare-metal server, it must be integrated with a Windows Deployment Services (WDS) server. You can enable this on the System Center Virtual Machine Manager host. However, in Chapter 2, you configured a virtual machine for this specific purpose.

This procedure walks through adding the WDS server to System Center Virtual Machine Manager. When this is completed, System Center Virtual Machine Manager uses the WDS

server as the engine for deployment and orchestrates the correct VHDX file deployment based on the physical computer profile you defined earlier in this chapter.

In a production environment, if a WDS-based PXE infrastructure is already configured, System Center Virtual Machine Manager can use it. However, in that case, System Center Virtual Machine Manager will initiate a bare-metal deployment to servers that it has designated as new VM or SOFS hosts. All other requests continue to be handled by the WDS server according to how it has been configured.

1. In the System Center Virtual Machine Manager console, click Fabric.
2. As shown in Figure 4-8, expand Infrastructure, right-click PXE Servers, and click Add PXE Server.




FIGURE 4-8 Adding a PXE server

3. In the Add PXE Server dialog box, in the Computer name box, type **WDS.contoso.com** as the name of the PXE server.
4. Specify an existing run-as account of SetupAdmin, and then click Add.
5. The Jobs window opens. Verify that the job has a status of Completed, and then close the window. The job sets up the new PXE server, installs the System Center Virtual Machine Manager agent on the PXE server, imports a new Windows Preinstallation Environment (Windows PE) image, and adds the machine account for the PXE server to System Center Virtual Machine Manager.
6. To verify that the PXE server is added, perform these steps:
 - A. Click Fabric, expand Servers, and then click PXE Servers.
 - B. On the Home tab, in the Show group, click Fabric Resources.
 - C. In the PXE Servers pane, verify that the PXE server appears with an agent status of Responding, as shown in Figure 4-9.

A screenshot of the 'PXE Servers (1)' pane in the System Center Virtual Machine Manager console. It displays a table with the following data:

Name	Agent Status	Role
WDS.contoso.com	Responding	PXE Server

FIGURE 4-9 A new PXE server

Procedure 9: Pre-provision Active Directory accounts

As part of the SOFS deployment, System Center Virtual Machine Manager works with Windows Server to automatically create appropriate DNS entries and Active Directory computer objects for the newly created operating systems. However, a best practice is to create DNS entries and Active Directory computer accounts in advance and allow time for DNS replication to occur. This step is not required, but it is strongly encouraged in an environment with multiple DNS servers where DNS replication might take some time.

1. Log on to DC01 using contoso\administrator credentials.
2. Open Active Directory Users And Computers, and in the navigation pane, expand contoso.com, and then click Computers.
3. Right-click Computers, click New, and then click Computer.
4. In the New Object - Computer text box, type **FS01**, and then click OK. When this is completed, right-click FS01, and select Disable Account. If prompted to confirm your choice, click Yes. The account must be disabled so that when the cluster is created, it can confirm that the account it will use for the cluster is not currently in use by an existing computer or cluster in the domain.
5. Repeat step 4 for FS02, FSCLUSTER, and SOFS.
6. Return to the main Active Directory Users And Computers window, click View, and select Advanced Features.
7. In the navigation pane, click Computers, right-click FSCLUSTER, and then click Properties.
8. In the FSCLUSTER Properties window, click Security.
9. Under Group Or User Name, scroll down and select Domain Admins (Contoso\Domain Admins). Note that all Domain Admin accounts have full permissions on this object. Because the System Center Virtual Machine Manager SetupAdmin account that you created uses contoso\administrator as the underlying account, SetupAdmin has the appropriate permissions on this particular object, along with the FS01 and FS02 computer objects. Click OK.
10. Right-click SOFS, and select Properties.
11. On the Security tab, click Add.
12. Click Object Types, make sure that Computers is selected, and then click OK. Then, under Enter The Object Name To Select, type **FSCLUSTER**, and then click OK. If a message appears, saying that you are about to add a disabled object, click OK.
13. Make sure that the cluster name account is selected and, next to Full Control, select the Allow check box. Click OK.
14. Close the Active Directory Users And Computers window, and open Domain Name System (DNS).
15. Under DC1, expand Forward Lookup Zones.

16. Right-click `contoso.com`, and create a new Host (A or AAAA) record. In the New Host window, type **FSCLUSTER** as the name. For IP address, enter **10.10.1.50**, select the check box to create an associated pointer record, and then click Add Host.
17. Repeat step 16, with **FSCLUSTER** and IP address **10.10.2.50**, click Add Host, and then click Done.

At this time, you will not specify DNS records for FS01 or FS02. At deployment time, System Center Virtual Machine Manager will dynamically allocate IP addresses (and thus configure DNS) to those new hosts from the static IP pools that you defined in Chapter 3, "Configuring network infrastructure." At that time, DNS records will be created for those machines.

18. Close the DNS window.

Procedure 10: Configure the WDS server with DHCP

Although System Center Virtual Machine Manager uses the BMC to wake the bare-metal physical servers, it still requires a DHCP server to be present. The DHCP server provides an IP address for one of the network adapters on the target bare-metal physical server. When an IP address has been provided, System Center Virtual Machine Manager can continue the configuration of the bare-metal server.

In this procedure, you add the DHCP role to the WDS server. Then you configure the appropriate settings.

1. Log on to your WDS VM, and enter the `contoso\administrator` credentials.
2. Open Server Manager. Click Manage, and then select Add Roles And Features.
3. On the Before You Begin page of the Add Roles And Features Wizard, click Next.
4. On the Select Installation Type page, select Role-Based Or Feature-Based Installation, and click Next.
5. On the Select Destination Server page, ensure that Select A Server From The Server Pool is selected, select `WDS.contoso.com` from the center pane, and click Next.
6. On the Select Server Roles page, select DHCP Server. When the Add Roles And Features dialog box opens, click Add Features, and then click Next.
7. On the Select Features page, click Next.
8. On the DHCP Server page, click Next.
9. On the Confirmation page, click Install. Installation can take a few minutes.
10. When installation is complete, click the notification icon, and then click Complete DHCP Configuration.
11. On the Description page, read the information, and then click Next.
12. On the Authorization page, ensure the Use The Following User's Credentials option is selected, and ensure that `contoso\administrator` is populating the User Name text box. Click Commit. This process can take a few minutes to complete.

13. When this action is completed, in Server Manager, click DHCP. In the center pane, right-click WDS, and select DHCP Manager. When it appears, expand the DHCP window.
14. In the left pane, expand DHCP, and then expand WDS.contoso.com.
15. Expand and then right-click IPv4, and select New Scope.
16. On the Welcome To The New Scope Wizard page, click Next.
17. On the Scope Name page, type **Deployment Scope** as the name, and then click Next.
18. On the IP Address Range page, type **10.10.0.50** as the Start IP address and **10.10.0.99** as the End IP address. Adjust the length of the subnet mask to 24, and click Next.
19. On the Add Exclusions And Delay page, click Next.
20. On the Lease Duration page, under Days, type **0**, under Hours, type **3**, and then click Next. Because you will use this only as a deployment scope, the lease can be short.
21. On the Configure DHCP Options page, select Yes, I Want To Configure These Options Now, and click Next.
22. On the Router (Default Gateway) page, type **10.10.0.254**, and click Add.
23. On the Domain Name And DNS Servers page, type **DC01**, and click Resolve. When the IP address populates, click Add.
24. Repeat step 23, but type **DC02** on the Domain Name And DNS Servers page, and then click Next.
25. On the WINS Servers page, click Next.
26. On the Activate Scope page, select the Yes, I Want To Activate This Scope Now option, and click Next.
27. On the Completing The New Scope Wizard page, click Finish.
28. Return to Server Manager, click WDS, and in the center pane, right-click WDS, and select Windows Deployment Services Management Console.
29. In the Windows Deployment Services console, in the top-left pane, expand Servers.
30. Right-click WDS.contoso.com, and select Properties.
31. In the WDS Properties window, click the DHCP tab.
32. Because DHCP is now running on this server, select both of the check boxes, and click OK. Close the Windows Deployment Services window, and return to DHCP.
33. In the DHCP window, expand Scope [10.10.0.0] Deployment Scope, and click Scope Options. You will see an option configured for WDS, labeled 060 PXEClient, as shown in Figure 4-10.

Option Name	Vendor	Value
003 Router	Standard	10.10.0.254
006 DNS Servers	Standard	10.10.0.11, 10.10.0.12
015 DNS Domain Name	Standard	contoso.com
060 PXEClient	Standard	PXEClient

FIGURE 4-10 New DHCP scope options

34. Click Server Options. Notice 060 PXEClient is listed there, also. Close the DHCP window.

Procedure 11: Discover and provision the SOFS with System Center Virtual Machine Manager

With the physical computer profile defined, the VHDX ready in the System Center Virtual Machine Manager temporary library, the WDS server integrated, and the relevant Active Directory objects pre-created, you can use System Center Virtual Machine Manager for the actual deployment of the SOFS.

When you create a SOFS cluster from bare-metal computers, the Create Clustered File Server Wizard does the following:

- Discovers the physical computers through out-of-band management
- Deploys the Windows Server 2012 R2 operating system image on the computers by using the physical computer profile (if configured to do so)
- Enables the file server role on the computers
- Enables the SOFS role on the cluster
- Adds the provisioned computers as a SOFS cluster under System Center Virtual Machine Manager management

The following steps walk you through the deployment of the SOFS to your bare-metal machines.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, and then click Fabric.
3. In the Fabric pane, click Servers. On the Home tab, in the Create group, click File Server Cluster. The Create Clustered File Server Wizard opens. On the General page, shown in Figure 4-11, do the following:
 - A. Type **FSCLUSTER** as the cluster name.
 - B. Type **SOFS** as the file server name.
 - C. Type **10.10.1.50** and **10.10.2.50** as the IP addresses for the cluster.
 - D. Click Next.

Enter a name for the cluster and for the file server on the cluster

Cluster name:

File server name:

If DHCP is not available, you must specify one or more IP addresses for the file server cluster. Use commas or new lines to separate multiple IP addresses.

Cluster IP addresses:

FIGURE 4-11 File server cluster deployment options

4. You are adding one IP address from each of the DataCenter_LN logical network static IP address pools for redundancy. You use a host address of 50 because it is outside of the static IP address pool range for the DataCenter_LN logical network. This prevents the cluster management IP addresses from conflicting with IP addresses provisioned to services or VMs on the DataCenter_LN logical network.
5. On the Provisioning Type page, shown in Figure 4-12, select the option to provision bare-metal computers with a new operating system by using a file server profile. Select the File Server physical computer profile, and then click Next.

Select the provisioning type for this cluster

Use existing servers running Windows Server 2012 R2
These servers must be in the same domain and must use the same administrator Run As account. They cannot have the Hyper-V role installed.

Run As account:

Provision bare-metal computers with a new operating system by using a file server profile
After deployment, the servers will be clustered. The selected profile provides the administrator Run As account and the domain name.

File server profile:

Skip cluster validation

FIGURE 4-12 Selecting the physical computer profile during file server deployment

6. On the Credentials And Protocol page, next to the Run As Account text box, click Browse, select SetupAdmin to access the BMC, and then click OK.
7. As Figure 4-13 shows, in the Protocol list, click Intelligent Platform Management Interface (IPMI) on port 623 for discovery, and then click Next. This selection can vary depending on your hardware.

Specify the account and protocol to be used for discovery

Specify the Run As account to use to discover computers that contain baseboard management controller (BMC) technology.

Run As account:

Select an available out of band management protocol to use for discovery.

Protocol:

Port:

FIGURE 4-13 Selecting the protocol and port used for bare-metal deployment of the SOFS cluster with System Center Virtual Machine Manager

8. On the Discovery Scope page, depending on how your BMCs were configured, select either the IP subnet or, more specifically, IP range. Whichever you choose, System Center Virtual Machine Manager scans the infrastructure and returns a list of discovered hosts. Hosts do not have to be powered on to be discovered. Click Next to start discovery.
9. When the hosts are discovered, on the Target Resources page, select the servers that you want to transform into the SOFS. In this case, select two servers. At this point, System Center Virtual Machine Manager uses the BMC to power on the servers. When the servers attempt to PXE boot, they are provisioned with a Windows PE image containing the System Center Virtual Machine Manager agent, which then performs a deep discovery of the physical hardware of the servers. This allows further customization of the deployment. Click Next.
10. On the Deployment Customization page, notice that deep discovery is running. This takes a few minutes. When deep discovery is completed, the servers are returned to a powered-off state. With the process complete, several warnings need to be corrected before you proceed. Select the first server in the list by selecting the first BMC IP address that has the warning triangle.
11. In the Computer Name text box, type **FS01**, and then click Network Adapters.
12. Two network adapters are listed, each with a different MAC address. Ensure that each of the network adapters shows Static IP under IP assignment. In your configuration, you might see more network adapters, depending on your hardware.
13. On the right side of the window, for the first network adapter, click the ellipses (...) to open the Properties window.
14. In the Network Adapter IP Configuration window, ensure that Specify Static IP Settings For This Network Adapter is selected.
15. In the Logical Network list, click the DataCenter_LN logical network. The default logical network is defined in the physical computer profile.
16. In the IP Subnet list, select the 10.10.1.0/24 subnet. The list of subnets is scoped to what is defined for the logical network in the associated network sites.

17. Ensure that the Obtain An IP Address Corresponding To The Selected Subnet option is selected, and then click OK.
18. Repeat steps 13 through 17 of this procedure to configure the second network adapter in this physical server. This time, use the 10.10.2.0/24 subnet.
19. Return to the Deployment Customization page, and for the first BMC IP address in the list, click Disks. If you want to select a specific disk for System Center Virtual Machine Manager to use for the operating system deployment, select the check box and then select the disk from the drop-down list . Otherwise, accept the default.
20. Repeat steps 10 through 19 of this procedure to configure FS02.
21. Return to the Deployment Customization page. The warnings should have disappeared, and you can proceed. Click Next.
22. On the Summary page, confirm the settings, and then click Finish to deploy the file server computers and to bring them under System Center Virtual Machine Manager management.
23. The Jobs window appears. This process takes some time to complete. The length of time depends on your environment. Ensure that all steps in the job have a status of Completed, and then close the window.
24. To confirm that the cluster was added, click Fabric, expand Storage, and then expand File Servers.
25. Verify that the name of the new SOFS cluster appears in the File Servers pane, and that its status is OK, as shown in Figure 4-14.

The screenshot shows a window titled "File Servers (1), File Shares (0)". Below the title bar is a search box. A table lists the file servers with the following columns: Name, Type, Cla..., Status, Total Capac..., Available..., Path, and Managed. One entry is visible: "SOFS.contoso.com" with Type "Scale-Out File Server", Status "OK", Total Capacity "0 GB", Available Capacity "0 GB", and Managed "Yes".

Name	Type	Cla...	Status	Total Capac...	Available...	Path	Managed
SOFS.contoso.com	Scale-Out File Server		OK	0 GB	0 GB		Yes

FIGURE 4-14 Fully deployed SOFS

You've successfully deployed a SOFS from bare metal using System Center Virtual Machine Manager.

Procedure 12: Check the cluster validation report

Whenever you construct a failover cluster, it is incredibly important to run a validation of the configuration. Cluster validation is an automated process that runs a set of focused tests on a collection of servers, networks, and associated storage that are planned for use as a failover cluster. The cluster validation process tests the underlying hardware and software to obtain an accurate assessment of how well failover clustering can be supported in a given configuration.

You can run validation on a particular configuration by using Failover Cluster Manager or Windows PowerShell. As another option—the option you use in this case—System Center Virtual Machine Manager will perform the task for you. The cluster validation process assesses

multiple areas, including cluster configuration, Hyper-V configuration, inventory, storage, and system configuration. To obtain Microsoft support, it is imperative that you run cluster validation for your configuration. As part of the deployment process, System Center Virtual Machine Manager ran a validation test before constructing the cluster.

1. Log on to FS01 using contoso\administrator credentials, and then open Failover Cluster Manager.
2. At the top left of the window, expand Failover Cluster Manager, right-click FSCLUSTER, and select View Validation Report.
3. Review the validation report to ensure there are no warnings or errors. Scroll through the different categories, and when you have reviewed the information, close the report, and then close Failover Cluster Manager.

TIP To learn more about cluster validation, visit TechNet at <http://technet.microsoft.com/en-us/library/jj134244.aspx>.

Procedure 13: Update software, drivers, and firmware

With your SOFS deployed, you can begin to aggregate the raw storage into pools. From there, you can create virtual disks that your workloads can use. Before you proceed, however, it's important to check and update, if required, the firmware for the key hardware of this storage configuration. If you have acquired new hardware, this procedure is less likely to be necessary. But if you are repurposing existing hardware, this procedure ensures that performance and reliability of the configuration is optimal.

For the cluster nodes, the VHDX you used for deployment should have been relatively up to date, meaning you shouldn't be missing too many, if any, Windows Server updates. However, checking it is straightforward, so complete the following steps.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open Cluster-Aware Updating. In the Connect To A Failover Cluster drop-down list, type **FSCLUSTER**, and click Connect. In this configuration, System Center Virtual Machine Manager is a cluster node itself. Therefore, System Center Virtual Machine Manager has the Cluster-Aware Updating interface available for use locally. You will use this to administer the FSCLUSTER remotely.
3. FS01 and FS02 should be displayed in the main window. On the right side, under Cluster Actions, select Analyze Cluster Updating Readiness. This checks that you can use Cluster-Aware Updating to update the cluster. You might receive warnings related to machine proxy and self-updating mode, but for this configuration, you can safely ignore them. Click Close.

NOTE Depending on your firewall configuration, you might receive an error during this step. If so, on both FS01 and FS02, run the following command from an administrative Windows PowerShell prompt:

```
Set-NetFirewallRule -Group "@firewallapi.dll,-36751" -Profile Domain -Enabled true to enable automatic restarts on a cluster node is updated
```

4. When the check is completed, under Cluster Actions, click Preview Updates For This Cluster.
5. In the Preview Updates window, click Generate Update Preview List. This might take a few minutes to complete. If no updates are listed, your nodes are up to date, and you can exit the Cluster-Aware Updating Wizard.
6. If updates are needed, close the Preview Updates window, and then click Apply Updates To This Cluster.
7. In the Cluster-Aware Updating Wizard, click Next.
8. On the Advanced Options page, review the defaults, and click Next.
9. On the Additional Options page, select the Give Me Recommended Updates The Same Way That I Receive Important Updates check box, and click Next.
10. On the Confirmation page, review the summarized information, and click Update. The wizard starts the process. You can safely click Close to return to the main window.

Cluster-Aware Updating triggers update scans on each node and orchestrates the updating of each node while ensuring the SOFS role, which is running on top of this cluster, is always readily available. Depending on how many updates each node requires, this process might take a few minutes.

With both nodes up to date from a Windows Server perspective, you should ensure that all firmware and drivers for the key hardware are up to date. For this configuration, for most hardware you are using drivers that are included with Windows Server and supplemented through Windows Update. However, specifically for the network cards, additional functionality might only be accessible when you're installing the hardware vendor's latest drivers and utilities. An example of such additional functionality is Datacenter Bridging (DCB), which you enable in the next chapter.

To reduce impact on the SOFS configuration, stagger the driver deployment across the two cluster nodes.

NOTE Refer to your hardware vendor's guidance to update firmware and, subsequently, drivers for key hardware. Updates should include, but not be limited to, server BIOS, SAS HBAs, network adapters, JBOD enclosure firmware, and JBOD drive firmware.

1. Log on to FS01 using contoso\administrator credentials.
2. Open Failover Cluster Manager. In the top-left side of the window, expand Failover Cluster Manager, and then click Nodes.
3. Right-click FS02, click Pause, and then click Drain Roles. This ensures that you can perform any manual updates and maintenance on FS02 without impacting the cluster.
4. Right-click FS02 again, and select Remote Desktop.
5. When you're logged in to FS02, apply any required firmware and driver updates and reboot as necessary. When finished, return to FS01.
6. In Failover Cluster Manager on FS01, under Nodes, right-click FS02, select Resume, and then select Fail Roles Back.
7. Right-click FS01, click Pause, and then click Drain Roles.
8. Apply any required firmware and driver updates and reboot as necessary. When finished, return to Failover Cluster Manager on FS01.
9. In Failover Cluster Manager on FS01, under Nodes, right-click FS02, select Resume, and then click Fail Roles Back.

Your cluster nodes are now up to date. Although this part of the process is manual, if you have a small number of nodes, the process is relatively short.

TIP Cluster-Aware Updating is extensible and can be configured to deploy hotfixes, firmware, and drivers as part of the process. You can learn more about this on TechNet at <http://technet.microsoft.com/en-us/library/jj134213.aspx>.

Procedure 14: Create storage classifications

Storage classification is used to differentiate between storage types based on either performance or guarantees that the underlying storage device offers. System Center Virtual Machine Manager uses storage classification for VM placement (which is defined in a template or at creation time) and for scoping specific storage for self-service users consuming clouds.

In this procedure, you create two storage classifications: one to identify storage that will be used by tenant VMs, and one to identify storage that will be used by infrastructure workloads and not exposed to tenants.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, and then click Storage.
3. Right-click Classifications And Pools, and select Create Storage Classification.
4. For the name, type **Tenant Storage**, and click OK.
5. Repeat steps 3 and 4 for two new classifications, one named **Infrastructure Storage** and another named **Primary Pool**.

Procedure 15: Create a storage pool

With your SOFS deployed, you might be wondering if you're all set. Can you now deploy VMs and other important data onto this new, resilient storage back end? The answer, at this stage, is no, not just yet. You still need to perform several steps to aggregate the physical disks into a storage pool and slice this pool into virtual disks, which are sometimes known as storage spaces. From there, you need to transform these into Cluster Shared Volumes (CSVs), on which file shares will reside. That might seem like a significant number of steps, but fortunately, System Center Virtual Machine Manager makes it easy. The example in Figure 4-15 shows the relationship between the disks, the pools, the virtual disks, the CSVs, and, finally, the file shares themselves.

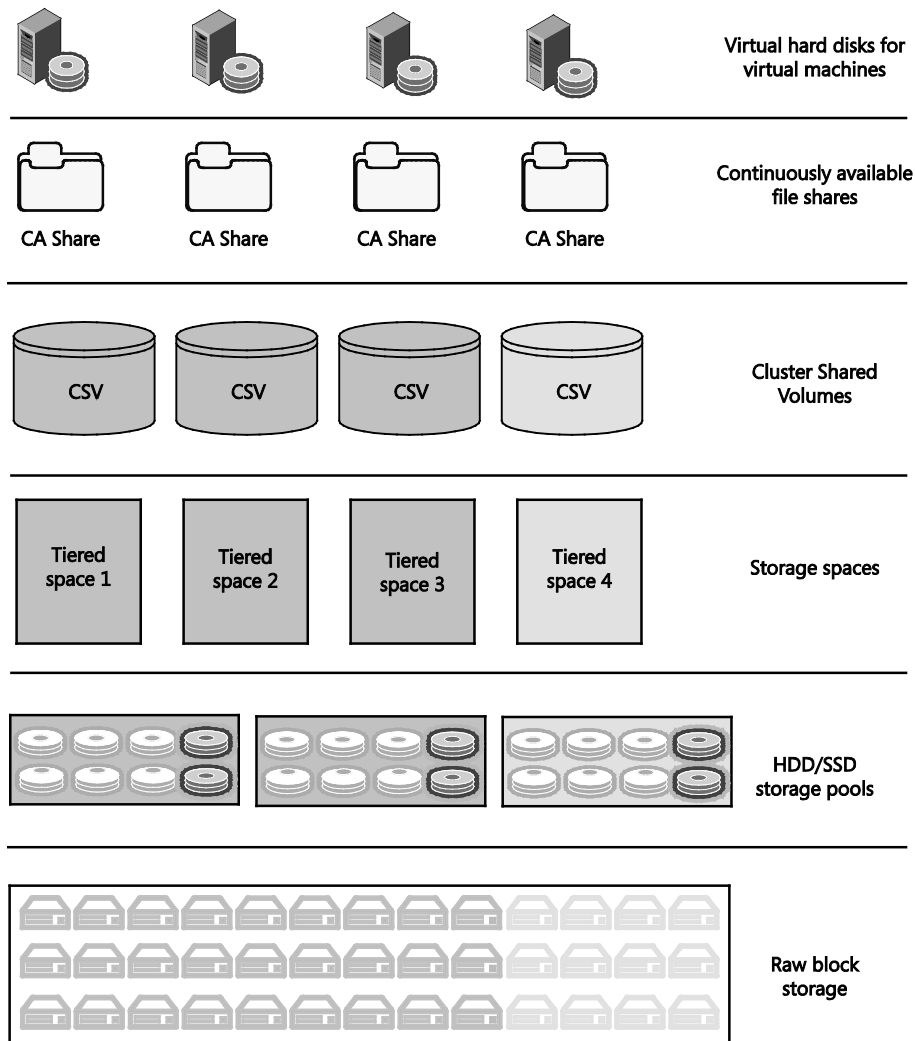


FIGURE 4-15 Layers of the Windows Server storage stack

As shown in the diagram, at the very bottom are the raw physical disks that are available through the JBODs. These physical disks are aggregated into one or more storage pools. Since this configuration uses multiple clustered hosts as part of a SOFS, these storage pools are also clustered, with each node having visibility into the pool.

From there, the total raw capacity is sliced into virtual disks, also known as storage spaces. These storage spaces can be configured to support tiering if the underlying storage pool has a mix of SSDs and HDDs available. You can also define the level of redundancy you require, with settings such as mirroring or parity.

Because this is a SOFS configuration, to ensure that these storage spaces are available across multiple nodes in the SOFS, they are transformed into CSVs. When you use CSVs with the SOFS, all cluster nodes can simultaneously write to the same storage, increasing performance and availability.

Finally, in this procedure you create continuously available file shares that sit within the CSVs. Continuously available file shares hosted on the SOFS let you store Hyper-V VM configuration files and virtual hard disks in easy-to-manage, remotely accessible file shares without sacrificing performance or availability.

Again, although this looks like a long process, System Center Virtual Machine Manager makes it straightforward.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, and then click Storage.
3. Click File Servers, and in the main window, right-click SOFS.contoso.com, and select Manage Pools.
4. In the Storage Pools window, click New.
5. In the Create Storage Pool window, for the name, type **POOL1**.
6. From the Classification drop-down list, select Primary Pool.
7. Within the Disk column, as shown in Figure 4-16, select all of the disks across both enclosures, and then click Create.

Disk	Interface	Media Type	Enclosure	Slot	Size	Status	Health
<input checked="" type="checkbox"/> PhysicalDisk6	SAS	HDD	0	5	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk19	SAS	HDD	1	7	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk21	SAS	HDD	1	9	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk13	SAS	HDD	0	11	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk25	SAS	HDD	1	11	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk8	SAS	HDD	0	8	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk15	SAS	HDD	1	5	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk12	SAS	HDD	0	10	931.51 GB	OK	Healthy
<input checked="" type="checkbox"/> PhysicalDisk7	SAS	HDD	0	4	931.51 GB	OK	Healthy

FIGURE 4-16 Disks discovered by System Center Virtual Machine Manager

As shown in Figure 4-16, the data in both the Enclosure and Slot columns is populated. If these columns are blank, ensure that you are using the latest JBOD enclosure firmware and disk firmware and that there are no disk errors.

To support deployments that require an added level of fault tolerance, storage spaces can associate data with a particular JBOD enclosure. This capability is known as enclosure awareness. With enclosure awareness, if one enclosure fails or goes offline, the data remains available in one or more alternative enclosures.

If you use enclosure awareness with storage spaces, your JBOD must support SCSI Enclosure Services (SES). With enclosure awareness in place, depending on your configuration, you can tolerate one or more failed enclosures, as shown in Table 4-2. To tolerate one failed enclosure with two-way mirrored spaces, you need three compatible storage enclosures. To tolerate two failed enclosures with three-way mirrored spaces, you need five compatible storage enclosures. This configuration has only two enclosures, so it cannot tolerate the failure of a complete enclosure.

TABLE 4-2 Storage space resiliency and SES options based on the number of JBODs

	TWO JBODS	THREE JBODS	FOUR JBODS
2-way mirror	1 disk	1 enclosure	1 enclosure
3-way mirror	2 disks	1 enclosure + 1 disk	1 enclosure + 1 disk
Dual parity	2 disks	2 disks	1 enclosure + 1 disk

As shown in Table 4-2, the configuration for this POC aligns with the Two JBODs column. When you configure Storage Spaces in future steps, the level of resiliency you select will determine how many disk losses can be tolerated. As shown in Figure 4-17, POOL1 has 24 disks and is the primary pool. The total disks include the data storage and protection.

File server pools:				
Name	Data Disks	Classification	Status	Health
POOL1	24	Primary Pool	OK	Healthy

FIGURE 4-17 Configured storage pools in System Center Virtual Machine Manager

- On the Storage Pools page, click OK, and System Center Virtual Machine Manager creates the storage pool. In this POC configuration, and for other smaller environments, a single pool eases administration while still providing high levels of redundancy and performance. But as the number of disks grows, it is a best practice to split those disks into separate pools to minimize the time required to fail over the storage pool to another node.

9. In the main System Center Virtual Machine Manager console window, under Storage, click Classifications And Pools. Notice that your three classifications are listed, one of which has an associated storage pool.

NOTE To view the clustered storage pool that System Center Virtual Machine Manager has created, log on to FS01, open Failover Cluster Manager, expand Storage, and click Pools. The new Storage Pool should be listed.

MORE INFO For further detailed guidance on how to optimally configure your storage pools based on your environment, review the Storage Spaces design information on TechNet at <http://go.microsoft.com/fwlink/?LinkID=517497>.

Procedure 16: Create a witness disk for FSCLUSTER

With the storage pool created, you can provision virtual disks, also known as storage spaces. Referring back to Figure 4-15, you've moved up the stack, having aggregated the raw storage into a single, more manageable pool.

The first disk that you create by using System Center Virtual Machine Manager is exclusively for FSCLUSTER to assist with achieving cluster quorum.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, click Storage, and then click File Servers.
3. In the main window, right-click SOFS.contoso.com, and select Properties.
4. On the General page, review the information. At the bottom, select Use Disk Witness For This File Server From The Specified Pool.
5. Ensure POOL1 is selected from the drop-down list, and click OK.

System Center Virtual Machine Manager automates the creation of the storage space from POOL1 in this case. From there, System Center Virtual Machine Manager handles the necessary formatting of the volume and attaches it to the FSCLUSTER, subsequently reconfiguring the cluster quorum to use a node and disk majority. This increases the reliability of the cluster.

MORE INFO For a deeper understanding of the failover clustering quorum configuration, read the following blog post from cluster MVP David Bermingham:
http://blogs.msdn.com/b/microsoft_press/archive/2014/04/28/from-the-mvps-understanding-the-windows-server-failover-cluster-quorum-in-windows-server-2012-r2.aspx.

Procedure 17: Create the virtual disks and file shares

With FSCLUSTER configured, you can use System Center Virtual Machine Manager to create a new storage space, or virtual disk, for actually storing data. When a storage space is created, the volume is formatted appropriately and converted to a CSV. A continuously available file share, onto which data can be placed, is provisioned on that CSV. The storage pool that was created earlier consisted of a mix of HDD- and SSD-based media. When creating a new virtual disk, you have the option to specify whether to enable storage tiering for that particular virtual disk, which will create a virtual disk of administrator-defined capacity, using a mix of the HDD and SSD devices. Unfortunately, in System Center 2012 R2 Virtual Machine Manager, you are unable to create tiered storage spaces. But you can still create a tiered storage space one of two ways. The first way is to use Failover Cluster Manager and create the virtual disk, and subsequently a file share, by using a GUI. The second option is to use Windows PowerShell to create a tiered space, which provides additional granularity over columns and interleave.

This example walks through the GUI to create this single storage space, and ultimately, file share.

1. Log on to FS01 using contoso\administrator credentials.
2. Open Failover Cluster Manager, expand Storage, and click Pools.
3. Select Cluster Pool 1. In the information pane, ensure that the pool name is POOL1.
4. On the right side of the window, click New Virtual Disk.
5. On the Before You Begin Page, click Next.
6. On the Select The Storage Pool page, click POOL1, and then click Next.
7. On the Specify The Virtual Disk Name page, type **TenantDisk1**, select Create Storage Tiers On This Virtual Disk, and then click Next.
8. On the Select The Storage Layout page, select Mirror, and click Next.
9. On the Configure The Resiliency Settings page, select Two-way Mirror, and click Next.
10. On the Specify The Size Of The Virtual Sisk page, shown in Figure 4-18, under Faster Tier (SSD), click Specify Size and type **124**; under Standard Tier (HDD), click Specify Size, and type **1295**; and then click Next. This creates a 1.39-TB virtual disk in this configuration.

Faster Tier (SSD)	Standard Tier (HDD)
Free space: 844 GB	Free space: 7,432 GB
<input checked="" type="radio"/> Specify size:	<input checked="" type="radio"/> Specify size:
<input type="text" value="124"/> GB	<input type="text" value="1295"/> GB

FIGURE 4-18 Tiered pool configuration during provisioning

11. On the Confirm Selections page, review the summarized information, and click Create. Windows Server creates the new virtual disk.

12. On the View Results page, ensure that the Create A Volume When This Wizard Closes check box is selected, and click Close.
13. On the Before You Begin page, click Next.
14. On the Select The Server And Disk Page, select SOFS, and click Next.
15. On the Specify The Size Of The Volume page, click Next.
16. On the Confirm Selections page, review the summarized information, and click Create. Windows Server creates and formats the volume, transforms it into a CSV, and updates all relevant information to enable usage by the SOFS.
17. When the process is completed, click Close.
18. Remaining within Failover Cluster Manager, under Storage, click Disks. Notice that the Cluster Virtual Disk has been created.
19. Repeat steps 4 through 18 to create three additional virtual disks with the names **TenantDisk2**, **InfraDisk1**, and **InfraDisk2**. They should all have the same capacities as previously used.
20. Repeat steps 4 through 18 for one final disk, which will host the System Center Virtual Machine Manager library. The size of this one can be smaller for this POC. Use **Library** as the name, and configure the virtual disk with 10-GB SSD and 500-GB HDD. The other settings can remain the same as previously used.

To view more information about the virtual disks that you have created, on FS01, in Failover Cluster Manager, expand Storage, and click Disks. Note the new cluster virtual disks that have been created. In the information window, under Resiliency, review that information about the resiliency mode chosen. In this case, the resiliency mode is mirror. Also note information about column count and interleave.

Besides offering resiliency to drive failures, storage spaces also offer increased performance by striping data across multiple disks. Storage spaces describe a stripe by specifying two parameters, NumberOfColumns and Interleave:

- A stripe represents one pass of data written to a storage space, with data written in multiple stripes (passes).
- Columns correlate to underlying physical disks across which one stripe of data for a storage space is written.
- Interleave represents the amount of data written to a single column per stripe.

The NumberOfColumns and Interleave parameters, accessible via Windows PowerShell and WMI, determine the width of the stripe ($\text{stripe_width} = \text{NumberOfColumns} * \text{Interleave}$). The stripe width determines how much data and parity (in the case of parity spaces) Storage Spaces writes across multiple disks to increase performance available to applications. In this configuration, the creation wizard sets the default number of columns and interleave based on this configuration.

TIP For further detailed guidance on how to optimally configure your virtual disks, including column count settings, based on your environment, review the Storage Spaces design information on TechNet at <http://go.microsoft.com/fwlink/?LinkID=517497>.

The final step is to create a new file share that resides on the new CSV. Although you could perform this task within Failover Cluster Manager, in this procedure you perform this step within System Center Virtual Machine Manager.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, click Storage, and then click Providers.
3. Right-click FSCLUSTER.contoso.com, and select Rescan. This makes the newly created CSV visible to System Center Virtual Machine Manager. This process might take a few minutes.
4. Once the process is completed, under Storage, click Classifications And Pools. Under the expanded POOL1, is a number of new Logical Units, representing the virtual disks just created.
5. On the top ribbon navigation, click Create File Share.
6. In the Create File Share Wizard, for the name, type **TenantShare1**.
7. From the Storage Type drop-down list, select Volume.
8. From the Volume drop-down list, select your appropriate CSV that represents TenantDisk1. If you are unsure about which CSV maps to which virtual disk, return to Failover Cluster Manager on FS01, click Storage, and then click Disks. When you select a particular virtual disk, the information pane displays the correct CSV. If you created your virtual disks in the same order as described above, your TenantDisk1 should map to C:\ClusterStorage\Volume2.
9. In the Classification drop-down list, select Tenant Storage, and then click Create.
10. Repeat steps 5 through 9 to create four additional file shares named **TenantShare2**, **InfraShare1**, **InfraShare2**, and **LibraryShare**. For InfraShare1, InfraShare2, and LibraryShare, use Infrastructure Storage as the classification. Again, use Failover Cluster Manager on FS01 to assist mapping the particular CSV to the underlying virtual disk.
11. When this is completed, in the Fabric view, under Storage, click File Servers. Expand SOFS.contoso.com to verify the new file shares are available.

Procedure 18: Create a library virtual machine

Earlier, when provisioning your SOFS nodes from bare metal, you added MGMT01 as a temporary library that System Center Virtual Machine Manager could use to store important files and folders relating to the bare-metal deployment. Now that the SOFS is constructed and you have file shares available for the placement of data, you can use one of the file shares,

specifically LibraryShare, to store important System Center Virtual Machine Manager-related artifacts. However, the challenge is an existing System Center Virtual Machine Manager-managed SOFS, and its underlying nodes, cannot be directly used as a library server. That scenario is unsupported. There is no problem using the LibraryShare itself as the location where data is stored for the library server, it is just that the SOFS nodes, specifically FS01 and FS02, cannot be library servers. You therefore need to configure an additional VM that runs on MGMT02 and is configured as the dedicated library server. In this procedure, you use System Center Virtual Machine Manager to attach the LibraryShare to this new library server VM, configuring all of the relevant permissions automatically.

To create the library server VM, follow these steps on MGMT02:

1. On MGMT02, open File Explorer, and navigate to \\MGMT01\Exports.
2. Right-click Sysprep.vhdx, and copy it. Navigate to the local D:\Exports folder on MGMT02, and paste the file to D:\VHDs. Change the name to **Library.vhdx**.
3. Open Hyper-V Manager, and in the top-right corner, click New, and then click Virtual Machine.
4. Create a new VM with the following settings:
 - **Name** LIBRARY
 - **Location** D:\VMs
 - **Generation** 1
 - **Memory** 2048 MB (Dynamic Memory unchecked)
 - **Networking** TenantNetwork vSwitch
 - **Virtual Hard Disk** Use an existing VHD and browse to D:\VHDs\Library.vhdx (If you have not set up a D:\ drive in your environment, then browse to C:\VHDs instead of D:\VHDs.)
5. Accept remaining defaults, and create the VM.
6. When the VM is created, select LIBRARY, and click Start. Double-click Library to open a console to the VM.
7. After a few minutes, the out-of-box experience (OOBE) begins. Select the appropriate country or region, app language, and keyboard layout, and then click Next.
8. Accept the license terms, and then enter your local administrator password. (Use a standard password for all the local and domain accounts you set up during this POC.) Click Finish.
9. Log in to the VM with your local administrator credentials. Server Manager opens automatically. Click Local Server, and then click the link next to the word Ethernet.

10. The Network Connections page opens. Right-click the network adapter, select Properties, select Internet Protocol Version 4 (TCP/IPv4), and then click Properties. Set the following options, and then click OK:
 - **IP Address** 10.10.0.19
 - **Subnet Mask** 255.255.255.0
 - **Default Gateway** 10.10.0.254
 - **Primary DNS** 10.10.0.11
 - **Secondary DNS** 10.10.0.12
11. Return to Server Manager, and under Local Server, click the link next to Computer Name.
12. Click Change, replace the existing name with **LIBRARY**, and then select Domain.
13. Type **contoso.com** for the domain name, and click OK.
14. When prompted for credentials, use contoso\administrator with your standard password.
15. Click OK, and then click Close. When prompted to restart, click Restart Now.
16. When you are back online, log in to the library VM as contoso\administrator, and open Windows Update. Run an update check against your WSUS infrastructure, install any required updates, and reboot as necessary.

With the new library VM set up and configured, you're ready to bring it under the management of System Center Virtual Machine Manager, and from there, attach the existing LibraryShare from the SOFS.

Procedure 19: Configure the library VM as the System Center Virtual Machine Manager library

In this procedure, you bring the new library VM under the management of System Center Virtual Machine Manager. The Library VM acts almost as a proxy between the System Center Virtual Machine Manager management server and the file share, LibraryShare, located on the SOFS, without the need for the SOFS nodes, FS01 and FS02, to actually become System Center Virtual Machine Manager library servers themselves.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, expand Infrastructure, and click Library Servers. MGMT01 is the only one listed.
3. On the top ribbon navigation, click Add Resources, and select Library Server.
4. On the Enter Credentials page, select Use An Existing Run As Account, and click Browse.
5. Select SetupAdmin, click OK, and then click Next.

6. On the Select Library Servers page, next to Computer Name, type **LIBRARY**, and click Add. This might take a few minutes. When done, click Next. Figure 4-19 shows the result.




FIGURE 4-19 The library VM added as a library server

7. On the Add Library Shares page, for now, accept that there are no shares, and click Next.
8. On the Summary page, review the settings, and click Add Library Servers. This process might take a few moments.
9. When this is completed, in the bottom-left navigation, click Library.
10. Expand Library Servers, right-click library.contoso.com, and select Add Library Shares.
11. In the Add Library Shares window, notice that the shares from the SOFS are now listed, even though they do not technically exist on the library VM. Select LibraryShare and the corresponding box in the Add Default Resources column. Adding the default resources adds the ApplicationFrameworks folder to the library share. Resources in the ApplicationFrameworks folder include x86 and x64 versions of the Server App-V Agent, Server App-V Sequencer, Windows PowerShell cmdlets for Server App-V, and the Microsoft Web Deployment tool. The folder also includes scripts that you can add to application profiles in service templates to install virtual applications and Web applications during service deployment. Click Next.
12. On the Summary page, review the settings, and click Add Library Shares. This process might take a few moments. This process provides the new library VM with the appropriate permissions to the SOFS library share and deploys the default resources to that share. The data technically resides on the SOFS share, but System Center Virtual Machine Manager uses the library VM as a conduit to access and use the data.
13. Expand Library Servers, expand library.contoso.com, and click the new LibraryShare that is listed there. Notice that several resources in the Application Frameworks folder have been deployed to the share. These are useful when you deploy virtual services but are out of scope for this configuration.

Configuring compute infrastructure

At this point in the POC deployment, the storage infrastructure is operational and the network is configured. The next step in the process is to deploy your compute nodes. For this POC configuration, you will be using Microsoft System Center Virtual Machine Manager to deploy, configure, and manage a four-node compute cluster that will use the software-defined storage infrastructure for its backend storage.

In Chapter 2, “Deploying the management cluster,” you manually deployed two Hyper-V hosts that ultimately became MGMT01 and MGMT02. These are currently running your management infrastructure virtual machines (VMs), including SQL Server, System Center Virtual Machine Manager, Active Directory Domain Services, and more. All of these VMs are currently running on local storage. In the event of an outage, this configuration doesn’t provide the highest levels of redundancy or performance.

Before you deploy your new Hyper-V hosts, you can perform a few steps on the existing infrastructure to increase the redundancy and performance of the overall solution. Your first steps will be to bring the management hosts into System Center Virtual Machine Manager’s management control. From there, you will connect them to the shared storage you constructed in Chapter 4, “Configuring storage infrastructure.” When they are attached, you’ll convert MGMT01 and MGMT02 into MGMTCLUS and migrate the storage of several workloads onto the Scale-Out File Server (SOFS) storage infrastructure. Figure 5-1 diagrams the management cluster, which uses tiered storage.

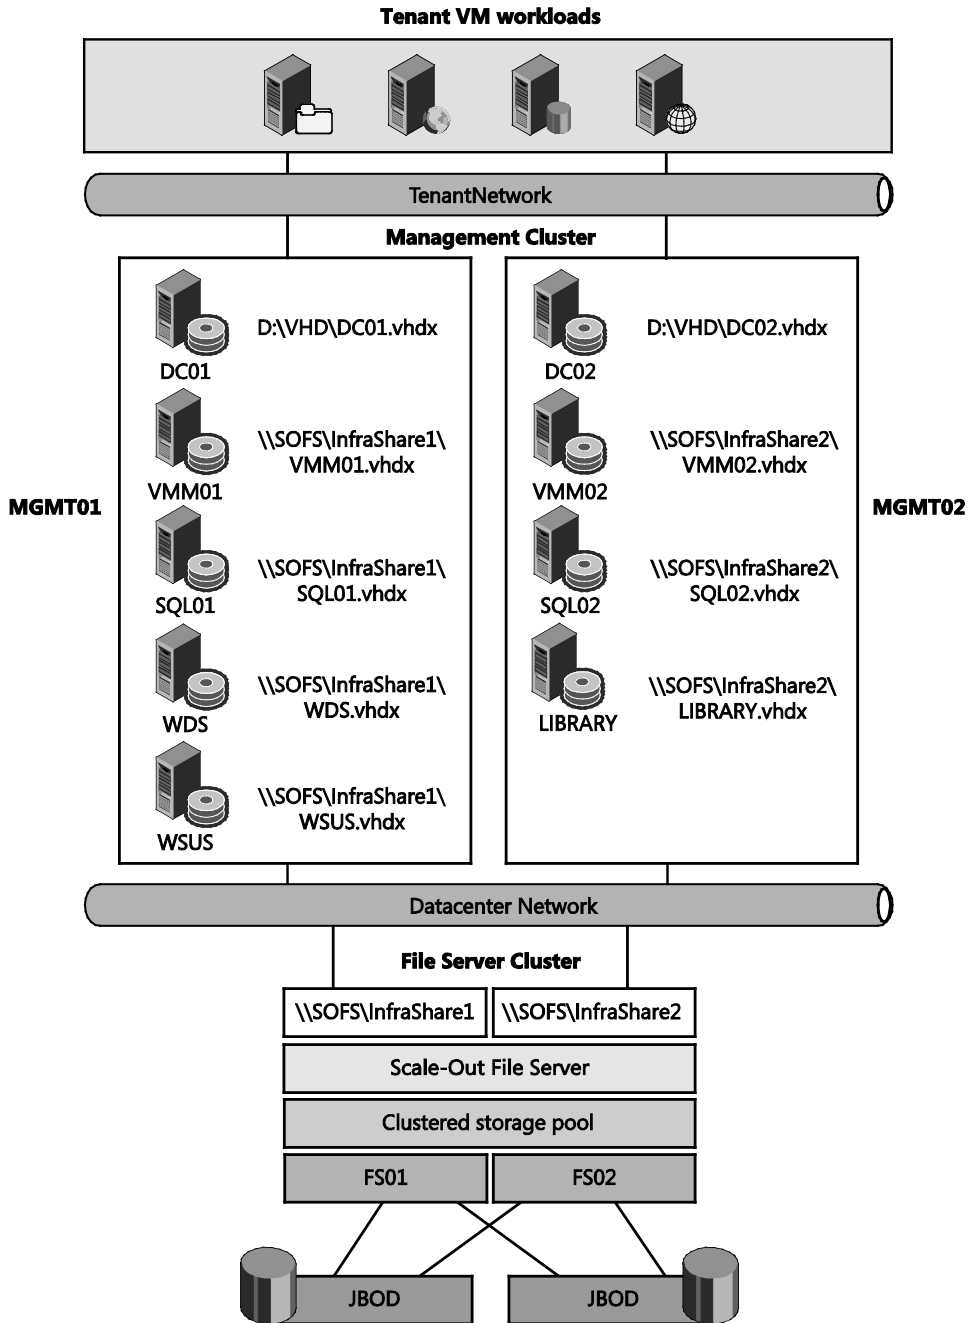


FIGURE 5-1 Logical view of management cluster using tiered storage

When that is complete, you'll begin the process of deploying the new Hyper-V hosts that will run your tenant VMs. As you did when you deployed the SOFS, you'll first construct a physical computer profile within System Center Virtual Machine Manager to standardize the process of building the Hyper-V hosts. You'll then walk through the process of bare-metal discovery and provisioning, deploying a Hyper-V image across to the bare-metal machines. When deployment is complete, you'll transform these hosts into a compute cluster, using the SOFS as the cluster's redundant, high-performance storage. Figure 5-2 diagrams the compute cluster, which uses SOFS storage.

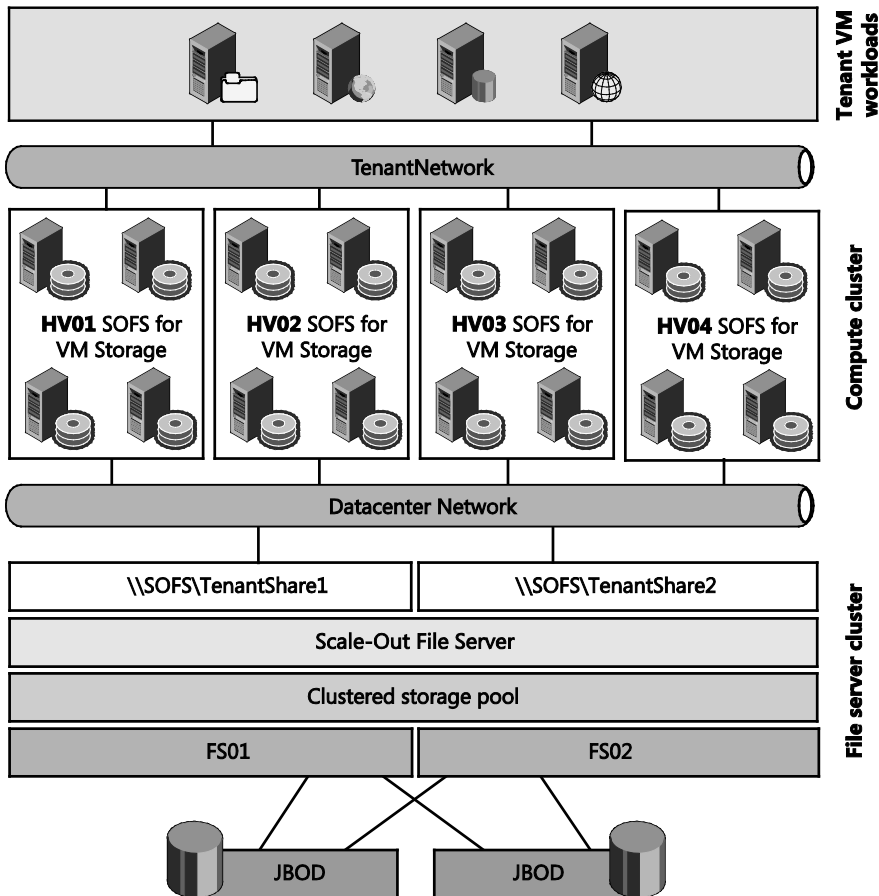


FIGURE 5-2 Logical view of compute cluster using SOFS storage

With the compute nodes deployed and configured, the final step involves configuring Data Center Bridging (DCB) and Remote Device Memory Access (RDMA) over Converged Ethernet (RoCE) to enable the optimization of traffic across the datacenter networks that you configured in Chapter 3, "Configuring network infrastructure."

Configuration walkthrough

In the following sections, you will walk through the key steps outlined in the previous section. You'll be taking advantage of the fact that you completed the storage and networking setup in the previous chapters. This existing setup will enable you to streamline the deployment of the Hyper-V nodes centrally from System Center Virtual Machine Manager and quickly enable a redundant configuration for tenant VMs.

NOTE The Hyper-V nodes will use four identically configured Dell PowerEdge R620 servers with the following specifications:

- 128-GB RAM, dual 2 Ghz Intel Xeon E5-2650, each with eight cores
- Two 10-Gbps RoCE-capable NICs, used for DatacenterNetwork 1 and DatacenterNetwork 2
- Four 1-Gbps NICs, used for TenantNetwork to allow VMs to communicate over the physical network
- One baseboard management controller (BMC) port for lights-out management over the network

Procedure 1: Create host groups in System Center Virtual Machine Manager

You can use host groups to group VM hosts in meaningful ways. Often, such groups are based on physical site location and resource allocation. At the host group level, you assign several settings and resources, such as custom placement rules, host reserve settings for placement, dynamic optimization and power optimization settings, network resource inheritance, host group storage allocation, and custom properties. By default, child host groups inherit the settings from the parent host group.

1. Log on to your VMM01 VM using contoso\administrator credentials.
2. From the Desktop, launch the System Center Virtual Machine Manager console. For the name, type **VMM-HA**, and click Connect. By entering VMM-HA, you'll be connecting to the highly available System Center Virtual Machine Manager configuration you constructed in Chapter 2.
3. In the System Center Virtual Machine Manager console, open the Fabric workspace.
4. In the Fabric pane, expand Servers, right-click All Hosts, and then click Create Host Group.
5. System Center Virtual Machine Manager automatically creates a new host group that is named New Host Group, and the host group name is highlighted.
6. Type **MGMT**, and then press Enter.
7. Repeat steps 4 through 6 of this procedure to create another host group and name it COMPUTE.

Procedure 2: Import management hosts into System Center Virtual Machine Manager

You created management hosts, MGMT01 and MGMT02, in Chapter 2. Now, with the networking configured and the storage deployed, it's time to add these two hosts to System Center Virtual Machine Manager.

1. Still logged into VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, then All Hosts, and expand MGMT.
3. Right-click MGMT, and select Add Hyper-V Hosts And Clusters, as shown in Figure 5-3.




FIGURE 5-3 Adding Hyper-V hosts and clusters

4. On the Resource Location page, select Windows Server computers in a trusted Active Directory domain, and click Next.
5. On the Credentials page, select the Use An Existing Run As Account option, and click Browse.
6. Choose SetupAdmin from the list of accounts, click OK, and then click Next.
7. On the Discovery Scope page, click Specify Windows Server Computers By Names, enter **MGMT01** and **MGMT02** on separate lines in the text box, and then click Next.
8. On the Target Resources page, wait for the discovery wizard to complete, select the two servers, and then click Next.
9. On the Host Settings page, ensure the selected host group is MGMT, and then click Next.
10. On the Summary page, review your settings, and click Finish.
11. The Jobs window opens. Observe the jobs through to completion, and then close the window.

System Center Virtual Machine Manager deploys an agent to each node and brings the two hosts into the MGMT host group. This process might take a few minutes. When it's completed, the two hosts are listed in the main window when you select the MGMT host group.

NOTE When the job completes, the result might show Completed w/Info because Multi-Path I/O is not enabled for known storage arrays on either host. For the purpose of this POC configuration, you can safely ignore this.

Procedure 3: Remove library server from MGMT01

In the previous chapter, you added MGMT01 as a temporary library server to centrally deploy your SOFS nodes. With the SOFS now deployed, you can release MGMT01 from its role as a library server. This enables MGMT01 to focus exclusively on being a Hyper-V host to run key workloads. With your new Scale-Out File Server, sofs.contoso.com, operational and file shares available to place data on, it makes sense to migrate the SYSPREP.vhdx file from MGMT01 to the new library share on library.contoso.com.

1. In the System Center Virtual Machine Manager console, open the Fabric workspace, and click Library Servers.
2. In the main window, right-click MGMT01.contoso.com, and select Remove. System Center Virtual Machine Manager removes MGMT01 from the list of library servers but leaves it in place as a Hyper-V host.
3. Click the Library workspace.
4. Expand library.contoso.com, right-click LibraryShare, and click Explore.
5. When File Explorer appears, right-click in the window, click New, and then click Folder. Name this folder **VHDs**.
6. In the address bar of File Explorer, type `\\MGMT01`, and click the Exports folder, click SYSPREP, and then click Virtual Hard Disks.
7. In the Virtual Hard Disks folder, right-click the SYSPREP.vhdx file, and select Copy.
8. Click Back to navigate back to the LibraryShare folder on library.contoso.com.
9. Click the VHDs folder, right-click in the space, and then click Paste.
10. When done, close File Explorer.
11. In the System Center Virtual Machine Manager console, right-click library.contoso.com, and click Refresh. Your new VHDs folder, along with the SYSPREP.vhdx file, should appear.
12. Right-click the SYSPREP.vhdx file, and select Properties.
13. In the SYSPREP.vhdx Properties window, change the name to **WS2012R2**.
14. From the Operating System drop-down list, select Windows Server 2012 R2 and the edition most relevant to the SYSPREP.vhdx file you created in Chapter 2.
15. From the Virtualization Platform drop-down list, select Microsoft Hyper-V, and then click OK.

Procedure 4: Construct a management cluster

In this procedure, you transform MGMT01 and MGMT02 into a management cluster. By constructing a management cluster from the two nodes, you increase the redundancy of the workloads running on top of the cluster. If a management node fails, the VMs running on that node will automatically fail over to the other node in the cluster.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. On the ribbon, click Create, and then select Hyper-V Cluster.
3. In the Create Cluster Wizard, on the General page, type **MGMTCLUS** for the cluster name.
4. Select the Use An Existing Run As Account option, and click Browse.
5. Choose SetupAdmin from the list of accounts, click OK, and then click Next.
6. On the Nodes page, from the Host Group drop-down list, select MGMT.
7. Under Available Hosts, you should see MGMT01 and MGMT02. Select both, click Add to move them to the MGMT host group, as shown in Figure 5-4, and then click Next.




FIGURE 5-4 Selecting the MGMT host group

8. On the IP Address page, under Network, click the boxes.
9. For the 10.10.1.0/24 network, under Static IP Pool, from the drop-down list, select Datacenter_LN_Pool1.
10. For the 10.10.2.0/24 network, under Static IP Pool, from the drop-down list, select Datacenter_LN_Pool2, and then click Next.
11. On the Storage page, click Next.
12. On the Virtual Switches page, because you created Hyper-V virtual switches for MGMT01 and MGMT02 in Chapter 2, you do not need to create any more in this wizard. Click Next.
13. On the Summary page, review your settings, and click Finish.

System Center Virtual Machine Manager automates the creation of the new failover cluster. This involves the installation of any roles and features, such as failover clustering and full cluster validation, and the creation of the cluster. This part of the process might take a few minutes, but you can monitor progress in the Jobs window.

Procedure 5: Check the cluster validation report

Cluster validation is an important step in the cluster creation process. Cluster validation ensures that the underlying configuration is optimally connected, configured, reliable, and robust. With your cluster created and managed by System Center Virtual Machine Manager, you can use the System Center Virtual Machine Manager GUI to access the cluster validation report instead of having to navigate to one of the nodes.

See also To learn more about cluster validation, visit TechNet at <http://technet.microsoft.com/en-us/library/jj134244.aspx>.

To access the cluster validation report, complete the following steps.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, then All Hosts, and then MGMT.
3. Right-click MGMTCLUS, and select Properties.
4. Click Status. On the Status tab, click the URL next to Report. The validation report opens.
5. Review the validation report. It will include a warning that you have no shared storage attached to the cluster. At this point, however, that is not an issue.

Procedure 6: Assign file share storage to MGMT01 and MGMT02

With the resilient Scale-Out File Server in place and MGMTCLUS up and running, you can move on to assigning storage to the cluster. The SOFS provides the shared storage for the hosts and ultimately delivers increased redundancy and performance for the management VMs.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, then All Hosts, and select MGMT.
3. Right-click MGMTCLUS, and select Properties.
4. Click the File Share Storage tab, and then click Add.
5. In the Add File Share window, from the drop-down list, select the InfraShare1 share, and then click OK.
6. Repeat steps 4 through 5 of this procedure to add InfraShare2, and then click OK.

When this is completed, you have shared storage presented to MGMT01 and MGMT02. System Center Virtual Machine Manager automatically modifies the shares to assign the necessary permissions to access the storage for the clustered hosts that are running Hyper-V.

Procedure 7: Configure a file share witness for MGMTCLUS, VMMCLUSTER, and SQLCLUSTER

Although System Center Virtual Machine Manager automates the deployment of the failover cluster, one piece that is yet to be completed is the cluster quorum. When you deployed the SOFS in the previous chapter, you used System Center Virtual Machine Manager to quickly create and assign a virtual disk from the storage pool and attach it directly to the file server cluster to provide the additional quorum vote. Unfortunately, that same automation is not

possible for the Hyper-V management cluster you have just constructed because you cannot present a virtual disk from the SOFS to the Hyper-V cluster. That process also doesn't apply to the SQLCLUSTER and VMMCLUSTER you created in Chapter 2. Instead, you need to present a file share for use as the witness, which will provide the additional quorum.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, click Fabric, click Storage, and then click File Servers.
3. On the ribbon, click Create File Share.
4. In the Create File Share window, next to Name, type **Witness**.
5. Next to Storage Type, from the drop-down list, select Storage Pool.
6. Next to Storage Pool, ensure that POOL1 is selected.
7. Next to Classification, from the drop-down, select Infrastructure Storage.
8. Next to Size (MB), type **3072**.
9. Next to Redundancy, from the drop-down list, select Two-way, and then click Create.
10. When the file share is created, right-click Witness, and select Properties.
11. In the Witness Properties window, clear the check box next to File Share Managed By Virtual Machine Manager, and then click OK.

This WITNESS file share will not be provisioned to Hyper-V hosts to store their VMs and workloads. Instead MGMTCLUS, SQLCLUSTER, VMMCLUSTER, and the future compute cluster will use this particular file share as the file share witness.

12. Log on to FS01 using contoso\administrator credentials.
13. Open Failover Cluster Manager, and click Roles.
14. Click SOFS, and at the bottom of the window, click Shares.
15. Right-click Witness, and select Properties.
16. In the Witness Properties window, click Permissions, and then click Customize Permissions.
17. Click the Share tab, and then click Add.
18. At the top of the Permission Entry For Witness window, click Select A Principal.
19. In the Select User, Computer, Service Account Or Group window, click Object Types, select Computers, and then click OK.
20. Return to the Select User, Computer, Service Account Or Group window, type **MGMTCLUS** as the name, and click OK.
21. Under Permissions, select Full Control, and then click OK.
22. Repeat steps 15 through 21 of this procedure to add permissions for SQLCLUSTER and VMMCLUSTER.
23. In the Advanced Security Settings For Witness window, click the Permissions tab.

24. Click Add. At the top of the Permission Entry For Witness window, click Select A Principal.
25. In the Select User, Computer, Service Account, Or Group window, click Object Types, select Computers, and then click OK.
26. Return to the Select User, Computer, Service Account, Or Group window, type **MGMTCLUS** as the name, and click OK.
27. Under Basic Permissions, select Full Control, and then click OK.
28. Repeat steps 23 through 27 of this procedure to add permissions for SQLCLUSTER and VMMCLUSTER.
29. In the Advanced Security Settings For Witness window, click OK.
30. Return to the Witness Properties window, click Apply, and then click OK.

The permissions are now set, and the management, System Center Virtual Machine Manager, and SQL Server cluster objects are assigned the correct permissions on the file share that will be used as the witness. Now you can proceed to assigning the share as the file share witness for the three clusters.

Procedure 8: Assign the file share witness to MGMTCLUS, VMMCLUSTER, and SQLCLUSTER

In this procedure, you assign the WITNESS file share that you've recently created, as the file share witness for the three clusters, MGMTCLUS, VMMCLUSTER, and SQLCLUSTER.

1. Log on to MGMT01 using contoso\administrator credentials.
2. Open Failover Cluster Manager, and expand MGMTCLUS.contoso.com.
3. Right-click MGMTCLUS.contoso.com, select More Actions, the Configure Cluster Quorum Settings, as shown in Figure 5-5.




FIGURE 5-5 Assigning the file share witness

4. In the Configure Cluster Quorum Wizard, on the Before You Begin page, click Next.
5. On the Select Quorum Configuration Option page, click the Advanced Quorum Configuration option, and then click Next.
6. On the Select Voting Configuration page, ensure All Nodes is selected, and then click Next.
7. On the Select Quorum Witness page, select Configure A File Share Witness, and then click Next.
8. On the Configure File Share Witness page, click Browse.
9. In the Browse For Shared Folders window, type **SOFS**, and click Show Shared Folders.

10. Six shared folders should appear in the window as shown in Figure 5-6. Select Witness, and then click OK.




FIGURE 5-6 The Shared Folders windows showing six shared folders

11. Return to the Configure File Share Witness page, and click Next.
12. On the Confirmation page, review your settings, and click Next. The quorum is configured to use the new file share as a witness disk. When this is complete, click Finish.
13. Under Failover Cluster Manager, click MGMTCLUS.contoso.com. In the main window under Cluster Core Resources, ensure that the File Share Witness is listed and has a status of Online.
14. Repeat steps 1 through 13 of this procedure on VMM01 (for VMMCLUSTER) and SQL01 (for SQLCLUSTER).

Procedure 9: Migrate management virtual machines to shared storage

With your management hosts set up as a cluster and the quorum optimally configured, you can now migrate the storage and other relevant VM data such as the configuration files onto the shared storage. This means that if a node fails, the VMs can continue to run on the other node in the cluster.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, expand All Hosts, and select MGMT.
3. Right-click MGMTCLUS, and select Refresh.
4. Open the VMs And Services workspace.
5. Under VMs And Services, expand All Hosts, expand MGMT, and click MGMTCLUS. If any of your VMs are missing, right-click each Hyper-V host individually and select Refresh Virtual Machines.
6. Select MGMT01 from the navigation pane on the left. On MGMT01, you should have DC01, VMM01, SQL01, WDS, and WSUS running. DC01 will remain on local storage.
7. Right-click WSUS, and select Migrate Virtual Machine.
8. The Migrate VM Wizard appears, starting on the Select Host page. Intelligent placement ranks each host based on the characteristics of the VM.
9. Click Make This VM Highly Available to force the VM to be deployed as a highly available

VM on the cluster and require that the VM's files are placed on shared storage.

10. Select MGMT01 as the target host. You do not need to move the running state of the VM. Instead, you will move the underlying VM virtual disks and configuration file. At the same time, this will make the VM highly available. The Live Migration Wizard allows you to accomplish this. Click Next. Figure 5-7 shows the dialog box where you can select from the possible destination hosts.




FIGURE 5-7 Dialog box for selecting a destination host for live migration

11. On the Select Path page, click Browse, and then select \\SOFS.contoso.com\InfraShare1.
12. Click the Automatically Place All VHDs With The Configuration option, and then click Next.
13. On the Summary page, review your settings, and click Move. The VM's virtual disks and configuration files will migrate to InfraShare1 on the SOFS, and the VM is made highly available on MGMTCLUS. Prior to this step, the VM was not configured for high availability because it was already running when the cluster was constructed. This is why you needed to perform this extra step.
14. Repeat steps 7 through 13 of this procedure for WDS, SQL01, and VMM01.

TIP By default, Hyper-V is configured with a maximum of two simultaneous live migrations between two hosts. To adjust this, right-click the Hyper-V host and select Properties. Under Migration Settings, adjust the maximums to suit your needs.

15. When all the VMs on MGMT01 are moved, proceed to MGMT02. Select MGMT02 from the list of hosts.
16. Repeat steps 7 through 13 of this procedure for SQL02, VMM02, and LIBRARY on MGMT02, selecting InfraShare2 as the target file share for storage.

Procedure 10: Enable Dynamic Optimization

With your management cluster configured and VMs residing on the shared storage, you can take advantage of additional System Center Virtual Machine Manager features to optimize the ongoing operation of the virtualized infrastructure. One of these capabilities is Dynamic Optimization. With Dynamic Optimization, System Center Virtual Machine Manager automatically assesses the current demand on the Hyper-V hosts on a configurable schedule and determines whether the environment would be more efficient if certain VMs were moved to other hosts.

You can configure Dynamic Optimization on a host group to migrate VMs within host clusters, and you can specify frequency and aggressiveness. Aggressiveness determines the amount of load imbalance that is required to initiate a migration during Dynamic Optimization. By default, VMs are migrated every 10 minutes with medium aggressiveness.

When configuring frequency and aggressiveness for Dynamic Optimization, you need to factor the resource cost of additional migrations against the advantages of balancing load among hosts in a host cluster. By default, a host group inherits Dynamic Optimization settings from its parent host group.

1. On VMM01, in the System Center Virtual Machine Manager console, open the VMs And Services workspace.
2. Expand All Hosts, right-click the MGMT host group, and select Properties.
3. Click the Dynamic Optimization tab, shown in Figure 5-8. Notice that this host group, by default, inherits its Dynamic Optimization settings from the parent All Hosts host group. Override this by clearing the Use Dynamic Optimization Settings From The Parent Host Group check box.




FIGURE 5-8 Configuring Dynamic Optimization

4. Adjust the Aggressiveness slider to High. In a production environment, you might choose to be more conservative and test the impact of each level of aggressiveness to determine which is right for you.
5. Optionally, you may select the Automatically Migrate Virtual Machines To Balance Load At This Frequency (Minutes) check box. By default, the optimizer runs only every 10 minutes anyway, but this setting allows you to override with a different value. For this configuration, set it at 5 minutes.
6. You can also adjust the default thresholds at which Dynamic Optimization operates. For this configuration, leave the defaults. Click OK.
7. Return to the main window, right-click MGMTCLUS, and select Optimize. This triggers an immediate Dynamic Optimization Assessment with suggestions, if applicable, for migrations.
8. In the Optimize Host Cluster window, shown in Figure 5-9, after the assessment has completed, click Optimize to run suggestions for migrations.

Virtual Machine	Source Host	Destination Host
WSUS	mgmt01	mgmt02
SQL02	mgmt01	mgmt02
SQL01	mgmt01	mgmt02

FIGURE 5-9 Optimizing the host cluster

By enabling Dynamic Optimization, you can then enable Power Optimization. To meet resource requirements, System Center Virtual Machine Manager uses Power Optimization to save energy by turning off hosts that are not needed within a host cluster and turning them on when they are needed again.

By default, System Center Virtual Machine Manager always performs Power Optimization when the feature is turned on. However, you can schedule the hours and days during the week to perform power optimization. For example, you might initially schedule power optimization only on weekends, when you anticipate low resource usage on your hosts. After observing the effects of power optimization in your environment, you might increase the hours. As part of this configuration, you will not be enabling Power Optimization.

See also For more information about Dynamic Optimization and Power Optimization, refer to the TechNet article at <http://technet.microsoft.com/en-us/library/gg675109.aspx>.

Procedure 11: Create availability sets for related virtual machines

With Dynamic Optimization enabled, System Center Virtual Machine Manager migrates VMs between hosts to ensure that the demands of each VM can always be met and that the overall cluster is operating as efficiently as possible. While analyzing the cluster and the VMs, System Center Virtual Machine Manager looks at demand across CPU, memory, disk, and so on.

But System Center Virtual Machine Manager is not taking into account the workload or application running inside the VM. You could therefore find yourself in a situation where both SQL Server VMs or both System Center Virtual Machine Manager VMs reside on the same host. If that host suffers an outage, the whole SQL Server or System Center Virtual Machine Manager cluster goes down. You therefore need to put Availability Sets in place to ensure such related VMs reside on separate hosts from one another. This is not applicable to the DC01 and DC02 in this POC since they reside on local storage on their respective hosts and thus are not considered for Dynamic Optimization.

1. On VMM01, in the System Center Virtual Machine Manager console, open the VMs And Services workspace.
2. Expand All Hosts, expand MGMT, and click the MGMTCLUS to display a list of all VMs on the cluster.

3. Right-click SQL01, and select Properties.
4. In the SQL01 Properties window, click the Hardware Configuration tab, scroll down and click Availability.
5. Under Availability Sets, click Manage Availability Sets.
6. In the Manage Availability Sets window, click Create.
7. In the Create Availability Set window, type **SQL Server Cluster**, and then click OK.
8. Return to the Manage Availability Sets window, and click OK.
9. Return to the SQL01 Properties window, and click OK.
10. Right-click SQL02, and select Properties.
11. In the SQL02 Properties window, click the Hardware Configuration tab, scroll down and click Availability.
12. Under Availability Sets, click Manage Availability Sets.
13. In the list of available properties, shown in Figure 5-10, select SQL Server Cluster, click Add, and then click OK.




FIGURE 5-10 Selecting SQL Server Cluster from the available properties

14. Return to the SQL02 Properties window, and click OK. You have created a SQL Server Cluster availability set, which System Center Virtual Machine Manager will use to ensure that the two VMs that are part of that availability set, are kept on different Hyper-V hosts, as much as possible.
15. Repeat steps 3 through 14 of this procedure for VMM01 and VMM02, creating a new availability set with the name SCVMM Cluster.

With the management cluster fully configured and the VMs redundant and optimized, you can move on to deployment of the main compute cluster. This will build on much of what you have learned in this chapter and in Chapter 4.

Procedure 12: Rack and connect Hyper-V compute nodes

In this procedure, you physically deploy the hardware into the environment. This might involve the acquisition of new hardware or repurposing of existing hardware. If you already have hardware racked and cabled in your environment, ensure that the networking is configured following the instructions in the following procedure. Name the four Hyper-V nodes HV01, HV02, HV03, and HV04.

1. Rack all physical servers that will become the Hyper-V hosts.
2. Connect the power. Connect the network, as described below. Note that at this point, you are simply plugging in the hardware. The network configuration will be automatically completed at deployment time.

The Hyper-V nodes use three separate networks as follows:

- DatacenterNetwork Port 1 (10 Gbps, not teamed)
- DatacenterNetwork Port 2 (10 Gbps, not teamed)
- TenantNetwork Team (teamed, consisting of four 1-Gbps adapters)

Plug all NICs on each node into physical network ports across separate switches as follows:

- DatacenterNetwork Port 1 on each node to a switch that is RoCE-capable
- DatacenterNetwork Port 2 on each node to another switch, also RoCE-capable
- Half of the ports for TenantNetwork Team to one switch and half to another switch for redundancy

Procedure 13: Configure BMCs

In this procedure, you'll deploy the Hyper-V nodes using System Center Virtual Machine Manager and take advantage of the bare-metal provisioning capabilities that you used to deploy the SOFS. A key enabler of this bare-metal deployment capability is the use of the baseboard management controller (BMC). The BMC enables out-of-band management and allows System Center Virtual Machine Manager to power on the physical server, run discovery scripts to inventory the physical server, control the network boot, and ultimately power off the physical server—all before the installed operating system starts. For this to work, your BMCs must be configured in advance to ensure System Center Virtual Machine Manager can communicate with them.

Table 5-1 shows the BMC settings for the Hyper-V nodes. Note that each hardware vendor's BMC configuration utility might differ from those of other vendors, so refer to your hardware vendor's guidance on how to configure your BMC. After the BMC is configured, make a note of the username, password, and IP address.

TABLE 5-1 BMC settings for Hyper-V nodes

PHYSICAL NIC	PURPOSE	NAME
HV01 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.3 SM: 255.255.255.0 DG: 10.10.0.254
HV02 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.4 SM: 255.255.255.0 DG: 10.10.0.254
HV03 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.5 SM: 255.255.255.0 DG: 10.10.0.254
HV04 - BMC Controller NIC	BMC interface for PXE Boot	IP: 10.10.0.6 SM: 255.255.255.0 DG: 10.10.0.254

For a successful bare-metal deployment, your Hyper-V nodes must support one of the following out-of-band management protocols:

- Intelligent Platform Management Interface (IPMI) version 1.5 or 2.0
- Data Center Management Interface (DCMI) version 1.0
- System Management Architecture for Server Hardware (SMASH) version 1.0 over WS-Management (WS-Man)

Whichever protocol you are using, ensure that you have the latest version of firmware for your BMC model.

Procedure 14: Configuring a BMC administrator in System Center Virtual Machine Manager (optional)

For System Center Virtual Machine Manager to successfully control the Hyper-V nodes through the BMCs, you need to provide System Center Virtual Machine Manager with the specific credentials for the BMC, for the purpose of bare-metal deployment. You can easily achieve this in System Center Virtual Machine Manager by specifying a new run-as account that contains the credentials required for the BMC.

NOTE If you are using the same credentials for the Hyper-V nodes' BMCs that you used for the SOFS BMCs, skip this step since you already have an appropriate run-as account configured. If you are using different credentials, proceed with this step.

To configure a BMC administrator, complete the following steps.

1. Log on to your VMM01 VM using contoso\administrator credentials.
2. From the desktop, launch the System Center Virtual Machine Manager console. For the name, type **VMM-HA**, and click Connect. By entering VMM-HA, you'll be connecting to the highly available System Center Virtual Machine Manager configuration you constructed in Chapter 2.
3. In the bottom-left corner of the System Center Virtual Machine Manager console, click Settings.
4. Expand Security, and select Run As Accounts. On the ribbon, click Create Run As Account.
5. In the Create Run As Account Wizard, in the Name box, type **Hyper-V BMC Administrator**. Optionally, add a description.
6. In the User Name box, enter the user name that is relevant for your BMC settings.
7. Below the user name, type a password and password confirmation in the corresponding boxes.
8. Unless your BMC account is within your domain, clear the Validate Domain Credentials box, and click OK.

You now have a run-as account that corresponds to the BMC on your Hyper-V nodes. If your Hyper-V nodes have BMC credentials that differ from one another, you will need repeat the steps to add additional run-as accounts, each with a unique name for easy identification.

Procedure 15: Create a physical computer profile

Physical computer profiles define the standardized characteristics of a physical server deployment, resulting in deployment of either a Hyper-V host or a SOFS. The concept of using a physical computer profile for physical server deployment is similar to the concept of using VM templates for VM deployment. The following procedure describes how to create a physical computer profile, specifically for a Hyper-V host rather than for a SOFS as covered in Chapter 4.

1. In the bottom-left corner of the System Center Virtual Machine Manager console, click Library.
2. Right-click Physical Computer Profile, and select Create Physical Computer Profile, as shown in Figure 5-11.




FIGURE 5-11 Creating a physical computer profile

- On the Profile Description page, name the profile **Hyper-V Host**, type a description, ensure that the VM Host option is selected, as shown in Figure 5-12, and then click Next.

FIGURE 5-12 Naming the profile

- On the OS Image page, shown in Figure 5-13, click Browse, and select the VHDX file that you published to the System Center Virtual Machine Manager library earlier. For this POC configuration, select the Do Not Convert The Virtual Hard Disk Type To Fixed During Deployment check box, and then click Next. In a production environment, you should clear the check box.

FIGURE 5-13 Specifying settings on the OS Image page

- On the Hardware Configuration page, under Management NIC, click IP Configuration.
- Select Allocate A Static IP From The Following Network, and select the Datacenter LN logical network from the drop-down list, as shown in Figure 5-14.

FIGURE 5-14 Allocating a static IP address

- At the top of the window, click Add, and then select Physical Network Adapter.
- Repeat steps 5 and 6 of this procedure for Physical NIC #1.
- At the top of the window, click Add, and then choose Physical Network Adapter.
- Under Physical NIC #2, click Connectivity Properties.
- Click Connect This Physical NIC To The Following Logical Switch. From the drop-down list, select the logical switch you created in Chapter 3. This will automatically select the uplink port profile that you created in Chapter 3.

12. Repeat steps 9 through 11 of this procedure, adding another three physical NICs. When this is complete, you should have a Management NIC and Physical NIC #1 through #5. Two of the NICs will represent the 10-Gbps DatacenterNetwork adapters, and the remaining four NICs will be transformed into a logical switch, allowing VM tenant traffic to flow onto the physical network.
13. Under Disk And Partitions, click OS, and review the settings. You will be using 100 percent of the available local disk for the deployment, but you can adjust if required.
14. Under Driver Options, click Driver Filter. You can use System Center Virtual Machine Manager to include hardware-specific drivers as part of the deployment process. This section of the wizard assumes you have imported the drivers into the System Center Virtual Machine Manager library. This will not be covered as part of this configuration; you will use the built-in Windows Server drivers for the time being. Click Next.
15. On the OS Configuration page, complete the following steps, and then click Next:
 - A. Type **contoso.com** as the domain to join, and select the SetupAdmin run-as account as the credentials for joining the domain.
 - B. For the admin password, type your standard password for the POC. This will be used to set the local administrator account password on the server.
 - C. For identity information, type your company information (optional).
 - D. Enter your product key if you are using one for the POC (optional).
 - E. Select your time zone.
 - F. Select an answer file if you're using one for any of your Windows settings (optional).
16. On the Host Settings page, leave the paths blank for now; you will attach the file shares after the hosts have been deployed. Click Finish.
17. Monitor the job for successful completion, and then close the Jobs window.


Procedure 16: Discover and provision the Hyper-V hosts with System Center Virtual Machine Manager

With the physical computer profile defined, the VHDX ready in the System Center Virtual Machine Manager library, and the Windows Deployment Server (WDS) integrated, you can use System Center Virtual Machine Manager to deploy the Hyper-V nodes. When you create new Hyper-V hosts from bare-metal computers, at a high level, the Add Resources Wizard does the following:

- Discovers the physical computer through out-of-band management
- Deploys an operating system image on the computer through the host profile or the physical computer profile
- Enables the Hyper-V role on the computers
- Brings the computer under System Center Virtual Machine Manager management as a managed Hyper-V host

In the following procedure, you deploy Hyper-V to your bare-metal machines.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, and then click Fabric.
3. In the Fabric pane, click Servers.
4. On the Home tab, in the Add group, click Add Resources, and then click Hyper-V Hosts And Clusters. The Add Resource Wizard opens.
5. On the Resource Location page, click Physical Computers To Be Provisioned As Virtual Machine Hosts, and then click Next.
6. On the Credentials And Protocol page, next to the Run As account box, click Browse, select your BMC Administrator account, and then click OK.
7. In the Protocol list, click IPMI On Port 623 For Discovery, and then click Next. This selection might vary depending on your hardware.
8. On the Discovery Scope page, depending on how your BMCs were configured, select either the IP subnet or, more specifically, the IP range. Whichever you choose, System Center Virtual Machine Manager scans the infrastructure and returns a list of discovered hosts. Hosts do not have to be powered on to be discovered. Click Next to start discovery.
9. When discovery is completed, on the Target Resources page, select the boxes next to the four servers that you want to deploy Hyper-V onto, and then click Next.
10. On the Provisioning Options page, in the Host Group list, select COMPUTE as the target location for the new Hyper-V hosts, and from the drop-down list, select your Hyper-V physical computer profile. Click Next.
11. The Deployment Customization page shows that deep discovery is running. This means that System Center Virtual Machine Manager used the BMC to wake the servers and is capturing details about their configuration, settings, and more. This will take a few minutes. When this is completed, the servers are returned to a powered-off state. With the process complete, you must address several warnings before proceeding. Select one of the servers in the list by selecting the BMC IP address with the warning triangle.
12. In the Computer Name text box, type **HV01**, as shown in Figure 5-15, and then click Network Adapters.



BMC IP address: 10.10.0.3
SMBIOS ID: 4c4c4544-0052-4210-8048-b1c04f445831
Serial number: 1RBHDX1
Manufacturer: DELL

Computer name: HV01
 Skip Active Directory check for this computer name

CPU: Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz
Core count: 16
Memory: 127.96 GB

FIGURE 5-15 Entering the computer name

13. You should see six network adapters listed, as shown in Figure 5-16. If the window is cramped with information, stretch it to the right by dragging the right edge of the window.


MAC Address	Name	Logical Swit...	IP Assignm...	Management NIC
00:0A:F7:5D:59:A0	SLOT 2 Port 1	None ▾	Static IP ▾	Yes ▾ ...
00:0A:F7:5D:59:A2	SLOT 2 Port 2	None ▾	Static IP ▾	No ▾ ...
E0:DB:55:20:A1:74	NIC1	Tenant_LN ▾	N/A ▾	No ▾ ...
E0:DB:55:20:A1:75	NIC2	Tenant_LN ▾	N/A ▾	No ▾ ...
E0:DB:55:20:A1:76	NIC3	Tenant_LN ▾	N/A ▾	No ▾ ...
E0:DB:55:20:A1:77	NIC4	Tenant_LN ▾	N/A ▾	No ▾ ...

FIGURE 5-16 Six network adapters listed

In this configuration, two of the network adapters are 10 Gbps and will be placed onto the DatacenterNetwork. The remaining four will be aggregated into a team and have a logical switch deployed, allowing future VMs to communicate out onto the network.

14. Ensure that one of your 10-Gbps network adapters under Management NIC has Yes selected. Also ensure that both of the 10-Gbps adapters have Static IP selected under IP Assignment.
15. For the first of your 10-Gbps adapters, click the ellipses (...) button to further customize the NIC settings. This opens the Network Adapter IP Configuration window.
16. In the Network Adapter IP Configuration window, click Specify Static IP Settings For This Network Adapter.
17. For logical network, from the drop-down list, select Datacenter_LN.
18. For IP subnet, from the drop-down list, select 10.10.1.0/24, and then click the Obtain An IP Address Corresponding To The Selected Subnet check box. As part of the deployment process, System Center Virtual Machine Manager assigns a static IP address from the pool you created in Chapter 3. Click OK.
19. Repeat steps 15 through 18 of this procedure, but this time, select the 10.10.2.0/24 subnet.
20. The four remaining 1-Gbps adapters in this configuration will be used exclusively by VMs. Therefore, as part of the deployment, specify that these network adapters should have a logical switch assigned to them. For each of the 1-Gbps network adapters, in the Logical Switch column, ensure they all have the Tenant_LN_LS selected.
21. For the first of your 1-Gbps adapters, click the ellipses (...) button. This opens the Network Adapter IP Configuration window.
22. In the Network Adapter IP Configuration window, click the Connect This Physical NIC To The Following Logical Switch check box. From the drop-down list, select Tenant_LN_LS if necessary. This is the logical switch that you constructed in Chapter 3.

23. Under Apply The Following Uplink Port Profile To This Physical NIC, from the drop-down list, select the uplink port profile, as shown in Figure 5-17. You created only a single profile in Chapter 3, so this is the only one available, as shown in Figure 5-17. Click OK.



The screenshot shows a configuration window with two sections. The first section is titled "Connect this physical NIC to the following logical switch:" and has a checked checkbox. Below it is a dropdown menu with "Tenant_LN_LS" selected. The second section is titled "Apply the following uplink port profile to this physical NIC:" and has a dropdown menu with "Tenant_LN_UPP_c830e941-cb8f-4407-b905-a4f83d7c9c11" selected.

FIGURE 5-17 Selecting the uplink port profile

24. Repeat steps 21 through 23 of this procedure for the remaining 1-Gbps network adapters.
25. Click Disks. If you have more than one local disk in the target server, you can use the check box and drop-down list to select the disk you'd like to deploy Hyper-V onto.
26. Repeat steps 11 through 25 of this procedure for HV02, HV03, and HV04, and then click Next.
27. On the Summary page, review your customization settings, and then click Finish.

System Center Virtual Machine Manager begins the deployment process, transferring a copy of the virtual hard disk, which is currently stored in the System Center Virtual Machine Manager library, to each of the Hyper-V hosts. The host is then configured to natively boot from the virtual hard disk. System Center Virtual Machine Manager proceeds with the automated setup of the Hyper-V host, including joining the domain, enabling the Hyper-V role, and deploying the System Center Virtual Machine Manager management agent. When this is completed, the four new Hyper-V nodes appear under System Center Virtual Machine Manager management within the COMPUTE host group.

Procedure 17: Update drivers and firmware on Hyper-V hosts

With your Hyper-V hosts deployed, it's important to check the firmware for the key hardware of the hosts and update if required. If you have acquired new hardware, it is less likely this step will be necessary. But if you are repurposing existing hardware, this step ensures that performance and reliability of the configuration is high. Your main focus should be on the BIOS and the BMC, along with the firmware and drivers for your network adapters—both the 1-Gbps and 10-Gbps adapters—inside each of the new Hyper-V hosts.

Refer to your hardware vendor's guidance for obtaining and applying any firmware updates. When these have been applied and the server has been rebooted, proceed with the next procedure and configure your Hyper-V cluster.

Procedure 18: Construct the Hyper-V cluster

In this procedure, you transform your four new Hyper-V nodes into a Hyper-V cluster. By constructing a cluster from the four nodes, you increase the redundancy of the workloads running on top of the cluster. If a Hyper-V node fails, the VMs running on that node automatically fail over to the other node in the cluster.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. On the ribbon, click Create, and then select Hyper-V Cluster.
3. In the Create Cluster Wizard, on the General page, type **HVCLUSTER** for the cluster name.
4. Select the Use An Existing Run As Account option, and click Browse.
5. Select SetupAdmin from the list of accounts, click OK, and then click Next.
6. On the Nodes page, from the Host Group drop-down list, select COMPUTE, as shown in Figure 5-18.




FIGURE 5-18 Selecting COMPUTE

7. Under Available Hosts, all four Hyper-V nodes should be listed. Select them all, click Add, and then click Next.
8. On the IP Address page, under Network, select each of the datacenter networks.
9. For the first network, under Static IP Pool, from the drop-down list, select the datacenter logical network pool. Repeat for the second network that is listed, and then click Next.
10. On the Storage page, click Next.
11. The Virtual Switches page should show that your logical switch is deployed to each host as part of the deployment process, so no further virtual switches are necessary. Click Next.
12. On the Summary page, review your settings, and click Finish.

System Center Virtual Machine Manager automates the creation of the new failover cluster. This involves the installation of any roles and features, such as failover clustering, full cluster validation, and the creation of the cluster. This part of the process might take a few minutes, but you can monitor progress in the jobs window.

Procedure 19: Check the cluster validation report

Cluster validation is an important step in the cluster creation process. It ensures that the underlying configuration is optimally connected, configured, reliable, and robust. With your Hyper-V cluster created and managed by System Center Virtual Machine Manager, you can use the System Center Virtual Machine Manager GUI to access the cluster validation report instead of having to navigate to one of the nodes.

See also To learn more about cluster validation, visit TechNet at <http://technet.microsoft.com/en-us/library/jj134244.aspx>.

To check the cluster validation report, complete the following steps.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, expand All Hosts, and then select COMPUTE.
3. Right-click HVCLUSTER, and select Properties.
4. Click Status. On the Status tab, click the URL next to Report. The validation report opens.
5. Review the validation report.

A warning appears indicating that you have no shared storage attached to the cluster. At this point, however, that is not an issue. You will be adding file share-based storage in the next procedure.

Procedure 20: Assign file share storage to Hyper-V cluster

Your Hyper-V cluster is now constructed and deployed. Before you can place workloads onto the cluster, you need to ensure it has shared storage to host the virtualized workloads. Using System Center Virtual Machine Manager to handle the storage assignment streamlines the process, in contrast to manually assigning file shares and configuring appropriate permissions. Also, when System Center Virtual Machine Manager assigns storage to the cluster, it automatically configures access for all the Hyper-V hosts within that cluster, again saving time.

To assign file share storage complete the following steps.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, expand All Hosts, and then select COMPUTE.
3. Right-click HVCLUSTER, and select Properties.
4. On the File Share Storage tab, click Add.
5. In the Add File Share window, shown in Figure 5-19, from the drop-down list, select TenantShare1, and then click OK.




FIGURE 5-19 Selecting TenantShare1

6. Repeat steps 4 and 5 of this procedure to add the TenantShare2, and then click OK to close the HVCLUSTER Properties window.

When this is completed, you have shared storage presented to all four Hyper-V nodes. System Center Virtual Machine Manager automatically modifies the shares to assign the necessary permissions for the clustered hosts that are running Hyper-V to access the storage.

Procedure 21: Configure the file share witness for HVCLUSTER

Earlier, you used the Witness file share located on the SOFS as the file share witness for managing the System Center Virtual Machine Manager and SQL Server clusters. In this procedure, you adjust the permissions of the file share once more to allow the new Hyper-V cluster to use the share as its file share witness.

This particular file share will not be provisioned to Hyper-V hosts to store their VMs and workloads. It will exclusively act as the file share witness for the cluster.

1. Log on to FS01 using contoso\administrator credentials.
2. Open Failover Cluster Manager, and click Roles.
3. Click SOFS, and at the bottom of the window, click Shares.
4. Right-click Witness, and select Properties.
5. In the Witness Properties window, click Permissions, and then click Customize Permissions.
6. On the Share tab, click Add.
7. At the top of the Permission Entry For Witness window, click Select A Principal.
8. In the Select User, Computer, Service Account, Or Group window, click Object Types, select Computers, and then click OK.
9. Return to the Select User, Computer, Service Account, Or Group window, type **HVCLUSTER** as the name, and click OK.
10. Under Permissions, select Full Control, and then click OK.
11. In the Advanced Security Settings For Witness window, click the Permissions tab.
12. Click Add. At the top of the Permission Entry For Witness window, click Select A Principal.

13. In the Select User, Computer, Service Account, Or Group window, click Object Types, select Computers, and click OK.
14. Return to the Select User, Computer, Service Account, Or Group window, type **HVCLUSTER** as the name, and click OK.
15. Under Basic Permissions, select Full Control, and then click OK.
16. In the Advanced Security Settings For Witness window, click OK.
17. Return to the Witness Properties window, click Apply, and then click OK.

The permissions are now set, and the HVCLUSTER cluster object is assigned the correct permissions on the file share that will be used as the witness. You can proceed to assigning the share as the file share witness for the cluster.

Procedure 22: Assign the file share witness to HVCLUSTER

In this procedure, you assign the file share that has just been modified as the file share witness for HVCLUSTER. This provides an additional level of redundancy for the Hyper-V cluster.

1. Log on to HV01 using contoso\administrator credentials.
2. Open Failover Cluster Manager, and expand HVCLUSTER.contoso.com.
3. Right-click HVCLUSTER.contoso.com, select More Actions, and then click Configure Cluster Quorum Settings, as shown in Figure 5-20.




FIGURE 5-20 Configuring the quorum settings

4. In the Configure Cluster Quorum Wizard, on the Before You Begin page, click Next.
5. On the Select Quorum Configuration Option page, click the Advanced Quorum Configuration option, and click Next.
6. On the Select Voting Configuration page, ensure All Nodes is selected, and click Next.
7. On the Select Quorum Witness page, select Configure A File Share Witness, and click Next.
8. On the Configure File Share Witness page, click Browse.
9. In the Browse For Shared Folders window, type **SOFS**, and click Show Shared Folders.
10. Four shared folders should appear in the window. Select Witness, and then click OK.
11. Return to the Configure File Share Witness page, and click Next.
12. On the Confirmation page, review your settings, and click Next. The quorum is configured to use the new file share as a witness disk. When this is complete, click Finish.
13. Under Failover Cluster Manager, click HVCLUSTER.contoso.com. In the main window under Cluster Core Resources, the File Share Witness is listed and should have a status of Online.

Procedure 23: Enable Dynamic Optimization and Power Optimization

As you learned earlier, System Center Virtual Machine Manager has a powerful capability to help balance resource usage across your Hyper-V clusters. This capability helps VMs receive the resources they need and ensures that no single host becomes a bottleneck for resources. In addition, with Dynamic Optimization enabled, System Center Virtual Machine Manager unlocks the opportunity to use Power Optimization.

Through Power Optimization, System Center Virtual Machine Manager saves energy by turning off hosts that are not needed to meet resource requirements within a host cluster. It turns the hosts back on when they are needed again.

By default, System Center Virtual Machine Manager always performs power optimization when the feature is turned on. However, you can schedule the hours and days during the week when power optimization is performed. For example, you might initially schedule power optimization only on weekends, when you anticipate low resource usage on your hosts. After observing the effects of power optimization in your environment, you might increase the hours. For this configuration, you will enable Power Optimization to operate on the weekend only.

1. On VMM01, in the System Center Virtual Machine Manager console, open the VMs And Services workspace.
2. Expand All Hosts, right-click the COMPUTE host group, and select Properties.
3. Click the Dynamic Optimization tab. Notice that by default, this host group inherits its Dynamic Optimization settings from the parent All Hosts host group. To override this setting, clear the Use Dynamic Optimization Settings From The Parent Host Group check box.
4. Adjust the Aggressiveness slider to High.
5. Optionally, you may select the Automatically Migrate Virtual Machines To Balance Load At This Frequency (Minutes) check box, shown in Figure 5-21. By default, the optimizer runs only every 10 minutes anyway, but this setting allows you to override with a different value. In this case, set it at 5 minutes.




FIGURE 5-21 Overriding the default Dynamic Optimization value

6. You can also adjust the default thresholds at which Dynamic Optimization operates. For this configuration, leave the defaults. Click OK.

- At the bottom of the window, select Enable Power Optimization, and then click Settings. The Customize Power Optimization Schedule window, shown in Figure 5-22, opens.

Thresholds			
Hosts will be considered for power optimization if they can be evacuated without causing any remaining nodes of the cluster to fall below the following thresholds.			
Resource	Amount	Unit	
CPU	40	%	
Memory	4096	MB	
Disk I/O	0	IOPS	
Network I/O	0	%	

FIGURE 5-22 The Customize Power Optimization Schedule window

- Read the important information under Thresholds. The text explains that hosts will be considered for power optimization only if powering down the node doesn't cause the remaining nodes in the cluster to have fewer resources remaining than the values in the boxes below the explanation. Figure 5-22 shows that power optimization will operate only if each of the remaining Hyper-V hosts has more than 40 percent CPU and 4 GB of memory remaining. These figures can be adjusted. However, for this POC configuration, accept the defaults.
- On the schedule grid, shown in Figure 5-23, click each box to change the color from blue to white. White indicates that Power Optimization will not run. For this POC configuration, configure power optimization to run only on the weekend. Thus, the squares corresponding to Monday to Friday should be white. For speed, you can drag your mouse across the squares to change them from blue to white.

	Midnight (AM)						Noon (PM)																		
	12	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11	12
Sun	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue
Mon	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White
Tue	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White
Wed	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White
Thu	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White
Fri	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White	White
Sat	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue

FIGURE 5-23 Filling out the scheduling grid

- When the schedule is set, click OK to apply the Power Optimization settings, and then click OK again to close the COMPUTE Properties window.

See also For more in-depth information about Dynamic Optimization and Power Optimization, refer to the TechNet article at <http://technet.microsoft.com/en-us/library/gg675109.aspx>.

You might be wondering how System Center Virtual Machine Manager wakes the powered down Hyper-V hosts. To enable this, you gave System Center Virtual Machine Manager the settings for each of the BMCs for your Hyper-V hosts before the deployment of Hyper-V to the bare-metal physical servers. System Center Virtual Machine Manager uses these settings to

discover the physical servers and initiate the process of deployment. System Center Virtual Machine Manager remembers these settings for both the IP and the BMC administrative accounts. On completion of the deployment process, System Center Virtual Machine Manager maintains that information in the host properties, ensuring that these hosts are instantly compatible with the power optimization capabilities.

Procedure 24: Configure RDMA over Converged Ethernet (RoCE) on the compute, management, and storage clusters

Before you step through configuring RoCE, it's important to understand the approach to converged infrastructures and how Remote Device Memory Access (RDMA) adds value. In traditional virtualized infrastructures, host servers typically ship with a large number of network adapters. In a Hyper-V environment, it's common to see servers with eight, ten, or even twelve 1-Gbps network adapters.

Here's how these are deployed: Using ten 1-Gbps adapters as an example, two of those are dedicated to server management, allowing administrators to manage the hosts and allowing the hosts to communicate with one another. These would be teamed for redundancy.

Two other adapters, again teamed, are used for live migration traffic. These are dedicated since live migration traffic typically saturates the available bandwidth to accelerate the process as much as possible. Separating this traffic onto dedicated network adapters helps to ensure this activity doesn't consume available bandwidth required for other networks such as management.

Two additional adapters are again teamed but this time dedicated to Cluster Shared Volume (CSVs), typically used when you have Hyper-V hosts that are using iSCSI or fibre channel SAN storage to host their VMs. This is particularly important in the event of redirected I/O. Redirected I/O typically occurs if a particular Hyper-V host loses its direct connection to the underlying storage, yet other nodes in the Hyper-V cluster, can still access the storage. The disconnected host will access the underlying storage *via* another host, with the CSV network as the transport. In addition, this network is typically used as a cluster heartbeat—although in reality, the cluster uses all networks as a heartbeat unless explicitly configured not to.

You also need to dedicate adapters for the VMs to communicate with the network. Obviously, this also needs redundancy, so a minimum of another two network adapters are required there. When teamed, these adapters have a virtual switch bound to the team, and from there, VMs are able to communicate onto the network.

So far, that accounts for eight network adapters. With two more optionally dedicated to storage, either iSCSI or SMB, there are ten adapters—just for a single host. Multiply that number by what you'd need for a large cluster, and the number of required ports, cables, and switches becomes significant. In addition, each of these networks should be on different subnets. All of this results in an administrative challenge. And one big assumption is that your physical server has enough PCI slots to accommodate all of these network adapters!

The question is: Why require so many? There are three reasons. First, the additional network adapters provide the respective bandwidth required, particularly for the live migration, CSV, storage, and VM networks. Second, having multiple adapters is important for redundancy, which is incredibly important in a production environment. The final reason is isolation. Having those networks providing separate channels for the different traffic types is very important, and historically, that could only be achieved with separate network adapters and teams.

To address some of these challenges, especially the sheer number of adapters, cables, and switches required, hardware vendors are shipping servers with a significantly reduced number of network adapters. In some cases, the number of network adapters might only be two. For performance, they will run at 10 Gbps, but how can you deploy all of those separated networks required for live migration, management, storage, VM traffic, and more and at the same time achieve the levels of isolation and redundancy required?

Fortunately, a few options are available. The first, shown in Figure 5-24, focuses on the ability of Hyper-V to support virtual network adapters for the host operating system. As shown in Figure 5-24 the Hyper-V host has only two 10-Gbps adapters. Using Windows PowerShell, or preferably System Center Virtual Machine Manager, you can create a Hyper-V virtual switch (standard or logical switch) and apply that virtual switch to the teamed 10-Gbps adapters. This allows VMs to communicate with the network. However, at this stage, this approach does not solve the challenge of having dedicated networks for the other functions such as live migration or management.

At this point, you need to create additional virtual network adapters, but not for the VMs. Instead, these are for the management host operating system.

As Figure 5-24 shows, five host virtual network adapters have been created. Two are dedicated to storage, either SMB or iSCSI; one is dedicated to live migration; one is dedicated to CSVs; and a final virtual network adapter is for host management.

Each of these virtual network adapters looks and feels like a physical network adapter to the host management operating system. You can open the Network and Sharing Center and see the adapter listed. You can open the properties of the adapter, set a static IP address, enable or disable certain adapter functions, assign a VLAN, and more, just as you can with a physical network adapter. The key element is that all of these host virtual network adapters are attached to the virtual switch, which is bound to the teamed physical 10-Gbps adapters. As a result, the virtual network adapters have 20 Gbps of bandwidth to share, as well as redundancy through the underlying team.

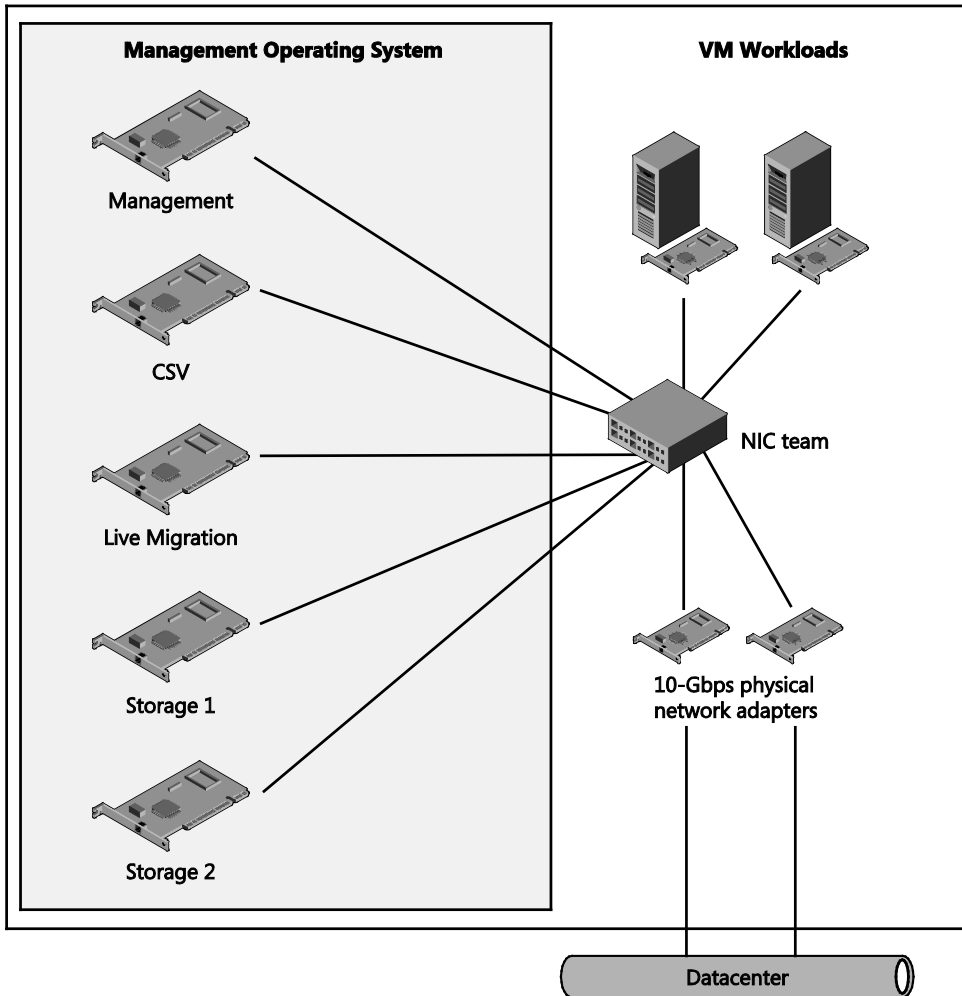


FIGURE 5-24 Logical view of virtual network adapters

To optimize this configuration, you can take advantage of specific Quality of Service (QoS) rules that you can apply to each of the host virtual network adapters. These QoS rules ensure that the virtual network adapters obtain the desired levels of bandwidth.

These rules are based on a weighting mechanism, but you can also configure them with absolute values. If you have 20 Gbps of available bandwidth to share across the host virtual network adapters and the VMs, you might decide to dedicate half of the available bandwidth to VMs. This gives you 50 percent of the remaining weight to distribute across the five host virtual network adapters.

By using the following Windows PowerShell command, you might configure, for example, the management virtual network adapter to have a guaranteed bandwidth of 5 percent:

```
Set-VMNetworkAdapter -ManagementOS -Name "Management" -MinimumBandwidthWeight 5
```


You might configure similar policies on the other host virtual network adapters. If you prefer, you can be more specific and specify a minimum and maximum bandwidth in bits per seconds rather than as weighting.

This example is just one of many possible converged combinations, and you'll find more useful examples for both converged and non-converged configurations on TechNet at <http://technet.microsoft.com/en-us/library/hh831441.aspx>.

This POC configuration has a slightly different configuration. However, now that you understand the basics of how to converge the networks into a reduced number of adapters, the POC configuration should be easier to comprehend. Note that as you continue, you will learn several new terms.

The key difference in the configuration that you have been working on in this guide is that the compute nodes and management nodes not only have the two 10-Gbps network adapters, but they also have four additional 1-Gbps network adapters. Figure 5-25 depicts the configuration.

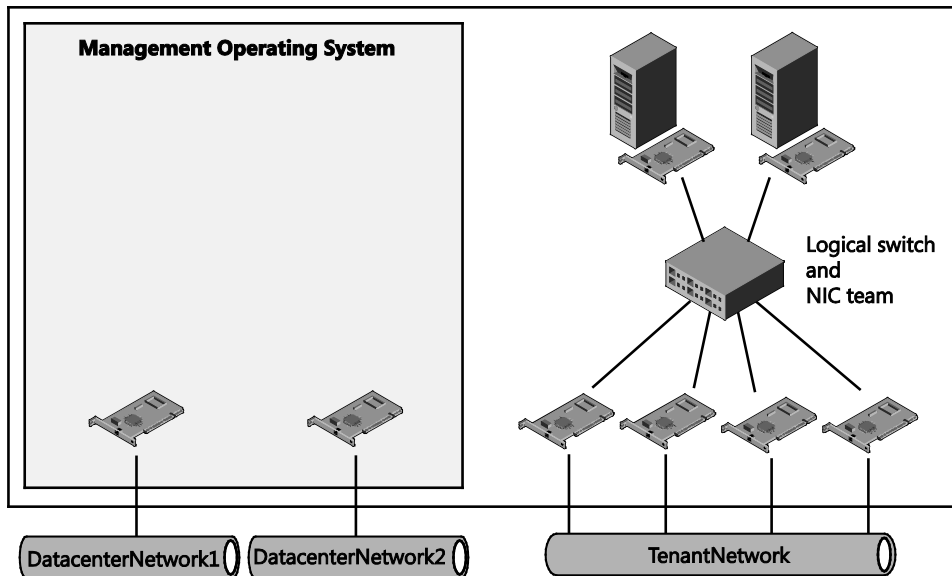


FIGURE 5-25 The configuration for this POC

As shown in the diagram, the management operating system is using the two 10-Gbps network adapters. The tenant VMs, because of the deployment of the logical switch, are sharing the four 1-Gbps network adapters. This provides the tenant VMs with redundancy thanks to the underlying NIC team and 4 Gbps of combined bandwidth.

One key difference between the 10-Gbps network adapters in this configuration and regular 10-Gbps network adapters is that the adapters for this POC support Remote Device Memory Access (RDMA). RDMA is, as the name suggests, the direct memory access from the memory of one machine to that of another machine without involvement of the operating

system on either machine. This enables higher throughput and low-latency networks and is highly advantageous for a virtualized infrastructure, where performance is key. The network adapter handles the transfer of data directly to or from application memory. Therefore, the operating system, in this case, Windows Server, does not need to be involved in the copy process. This frees up valuable CPU resource in the host and eliminates the requirement for any form of caching or context switches from the operating system side.

There are, however, many implementations of RDMA. Infiniband is one example. Although Infiniband is mature, it is typically on the expensive side, largely due to the bespoke nature of the adapters, cables, and switches. That said, Infiniband does support speeds of 54 Gbps and higher, which is desirable and will be supported well into the future. The major alternatives to Infiniband are Internet Wide Area RDMA Protocol (iWARP) and RDMA over Converged Ethernet (RoCE). Both use the more familiar TCP/IP protocol as transport, meaning they can work with many 10-Gbps switches that are already on the market. In both cases, the network cards handle the advanced processing and transfer, yet they can also be used just as regular network adapters.

The configuration for this POC uses RoCE, but the steps do not differ vastly if you have iWARP network adapters. If these network adapters are so powerful and optimized for high performance, you might be wondering, why doesn't this POC use a configuration like that shown in Figure 5-25. In Figure 5-25, the management operating has several host virtual network adapters attached to a virtual switch and bound to a team, which consists of the two 10-Gbps network adapters.

In fact, there are a number of reasons for the configuration in this POC. First, RDMA would not pass via a virtual switch. Therefore, the performance would be reduced, especially accessing the storage infrastructure over SMB 3.0. When you combine RDMA-capable network adapters with the SMB 3.0 protocol, a new capability known as SMB Direct is enabled. For storage environments, this provides a significant boost in performance. For example, your tenant VMs running on the compute nodes will be stored on the SOFS. The faster the access from the compute nodes over SMB 3.0 to the SOFS, the better the overall performance of the tenant workloads. SMB Direct is a key contributor in this case. By having multiple RoCE-capable network adapters, each on a separate datacenter network and thus subnet, you can combine SMB Direct with SMB Multichannel, increasing performance and redundancy still further.

Second, RDMA does not support teaming. Although you can team the network adapters, the RDMA capability would be disabled, and thus performance would be reduced. For these reasons, this POC uses the four 1-Gbps adapters for the tenant VMs, by means of the logical switch. The POC will dedicate the two RoCE-capable network adapters to other network needs, such as live migration, storage, management, and cluster communications.

The next question is if you can't create host virtual network adapters that are bound to the RoCE-capable network adapters, how do you distribute the 20 Gb of bandwidth across the five networks that are required (Host Management, Live Migration, CSVs, Storage 1, and Storage 2). To do that, you use a combination of Priority-based Flow Control (PFC) and DCB.

DCB is a collection of standards that defines a unified 802.3 Ethernet media interface, or fabric, for LAN and SAN technologies. DCB extends the current 802.1 bridging specification to support the coexistence of LAN-based and SAN-based applications over the same networking fabric within a datacenter. DCB also supports technologies such as Fibre Channel over Ethernet (FCoE) and iSCSI by defining link-level policies that prevent packet loss.

DCB consists of several 802.1 draft standards that specify how networking devices can interoperate within a unified datacenter fabric. The first is PFC. PFC supports the reliable delivery of data by substantially reducing packet loss resulting from congestion. This allows loss-sensitive protocols, such as FCoE, to coexist with traditional loss-insensitive protocols over the same unified fabric.

Similar to FCoE, RDMA is another loss-sensitive protocol. If RDMA is built on top of Ethernet directly, as it is with RoCE, the Ethernet transport must be lossless.

Traditionally, link-level flow control, which relies on the 802.3 Pause frame, is a solution. But link-level flow control causes problems, such as head of line blocking. PFC resolves this issue. Windows Server 2012 R2 allows you to enable PFC as long as the NIC supports it. When PFC is enabled for RoCE on both ends of an Ethernet link, only the virtual link designated for RoCE, which is denoted by a priority value, becomes lossless, and other workloads on the same physical link do not suffer from head of line blocking.

Alongside PFC, Enhanced Transmission Selection (ETS) is a transmission selection algorithm (TSA) that is specified in the IEEE 802.1Qaz draft standard. This standard is part of the framework for the DCB interface.

ETS allocates bandwidth between traffic classes that are assigned to different IEEE 802.1p priority levels. Each traffic class is allocated a percentage of available bandwidth on the data link between directly connected peers. If a traffic class doesn't use its allocated bandwidth, ETS allows other traffic classes to use the available bandwidth that the traffic class is not using.

Finally, the Data Center Bridging Exchange (DCBX) protocol allows DCB configuration parameters to be exchanged between two directly connected peers. This allows these peers to adapt and tune QoS parameters to optimize data transfer over the connection.

DCBX is also used to detect conflicting QoS parameter settings between the network adapter (local peer) and the remote peer. Based on the local and remote QoS parameter settings, the miniport driver resolves the conflicts and derives a set of operational QoS parameters. The network adapter uses these operational parameters for the prioritized transmission of packets to the remote peer.

NOTE The following steps assume the use of an RoCE-capable network adapter. If you are using iWARP network adapters, refer to the steps on TechNet, which, though similar, do differ slightly from those below (<http://technet.microsoft.com/en-us/library/dn583825>). If your network adapters are not RoCE-capable, yet still support DCB, you can proceed with the steps below. However, no SMB traffic will be offloaded directly to the network adapters.

To configure DCB and PFC for your environment, you'll use Windows PowerShell since the controls are not exposed through UI or through Windows PowerShell. You begin by configuring this on the file server cluster.

1. Log on to FS01 using `contoso\administrator` credentials.
2. Open a Windows PowerShell window as an administrator.
3. Run the following command to enable the DCB functionality, which is not enabled by default.

```
Install-WindowsFeature Data-Center-Bridging
```

4. To ensure you are configuring from a clean starting point, clear any existing PFC configuration that exist on this particular host.

Because System Center Virtual Machine Manager deployed your SOFS and underlying file server cluster nodes, it's unlikely that any PFC configurations are present. But it's worthwhile to apply this anyway, and it is useful if you make a mistake at any point.

To clear a previous configuration, run the following:

```
Remove-NetQoSTrafficClass  
Remove-NetQoSPolicy -Confirm:$False
```

5. Configure the DCBX protocol.

You can use the `Set-NetQoSDcbxSetting` cmdlet to configure network adapters in the server to accept DCB configurations from Windows Server or from a remote device via the DCBX protocol. If the `-Willing` parameter is set to `True`, Windows Server will not send PFC and traffic class settings to DCB-capable network adapters in the server. If the `-Willing` parameter is set to `False`, then Windows Server will send the settings to the network adapters. This POC configuration must be able to send the settings.

Therefore, configure the parameter as `False`.

```
Set-NetQoSDcbxSetting -Willing 0
```

6. Create the QoS policies and tag the different types of traffic that will operate over the RoCE-capable network adapters.

In this case, create four policies. The first is for SMB-based traffic, operating over port 445. It's important to note that in this configuration, the SMB protocol will be used for storage and live migration traffic, which brings a significant speed boost to both traffic types. The remaining policies are default, TCP, and UDP. TCP and UDP are created as separate policies from the default policy, with a different priority, because both the TCP and UDP protocols already have their own sophisticated algorithms to optimize flow control. Default is a catch-all policy for other traffic types that require flow control.

```

New-NetQoSPolicy "SMB" -NetDirectPortMatchCondition 445
-PriorityValue8021Action 3
New-NetQoSPolicy "DEFAULT" -Default -PriorityValue8021Action 3
New-NetQoSPolicy "TCP" -IPProtocolMatchCondition TCP
-PriorityValue8021Action 1
New-NetQoSPolicy "UDP" -IPProtocolMatchCondition UDP
-PriorityValue8021Action 1

```

7. With the policies created, enable PFC for a specific priority. In this configuration, priority 3 prioritizes the SMB-based traffic along with the default traffic, excluding TCP and UDP. Disable PFC for the other priorities.

```

Enable-NetQoSFlowControl -Priority 3
Disable-NetQoSFlowControl 0,1,2,4,5,6,7

```

8. Enable the policies and configuration on each of the RoCE-capable network adapters within this physical host. If you haven't already, you might want to rename your 10-Gbps network adapters to match the following example; otherwise, substitute your existing network adapter names.

```

Enable-NetAdapterQoS -InterfaceAlias "DatacenterNetwork Port 1"
Enable-NetAdapterQoS -InterfaceAlias "DatacenterNetwork Port 2"

```

9. The final step is optional. Across the two RoCE-capable network adapters, you have 20 Gb of bandwidth. This bandwidth will be consumed by SMB 3.0 storage traffic, along with bursts of consumption, through live migration over SMB. In addition, management tasks will also require bandwidth for tasks such as establishing a remote desktop connection to the server or for management agents to communicate to System Center Virtual Machine Manager. To control the amount of bandwidth a particular priority uses, establish a QoS policy for a particular traffic class.

```

New-NetQoSTrafficClass "SMB" -Priority 3 -BandwidthPercentage 80
-Algorithm ETS

```

This QoS policy focuses exclusively on priority 3 and uses the ETS algorithm to allocate bandwidth between traffic classes that are assigned different priority levels. In the above configuration procedure, by restricting this priority to 80 percent, both SMB and DEFAULT will be limited to consuming a maximum of 16 Gb from the 20 Gb total available bandwidth. That leaves 4 Gb of bandwidth for use by the TCP and UDP protocols.

10. Repeat steps 1 through 9 of this procedure on FS02, MGMT01, MGMT02, HV01, HV02, HV03, and HV04.

RoCE is now successfully configured on all of the key physical servers in the configuration. Before you can test, you need to ensure that you have configured the physical switch, or switches, that connect your physical servers.

Procedure 25: Configure RoCE on physical switches

For RoCE to perform reliably, you must enable PFC on the switch. The procedure for configuring PFC on switches varies depending on the type of switch you select. Refer to the hardware documentation for your switch for details on how to perform this step.

This is not an optional step. If you do not configure PFC on the switch, your RoCE-capable network adapters will perform poorly.

Procedure 26: Test RoCE configuration and connectivity

With your switches and the network adapters on each of the key nodes configured, you can test the connectivity. Specifically, in this procedure you test the use of RoCE (or iWARP) to confirm that all is operating as expected.

1. Log on to MGMT01 using contoso\administrator credentials.
2. Open a Windows PowerShell window as an administrator.
3. Use the Net TCP/IP cmdlets in Windows PowerShell to verify that Network Direct is globally enabled and that you have RDMA-capable network adapters. In Windows PowerShell, type the following on both the SMB server and the SMB client. Run the following three commands, one after another. The first command should return "Enabled." The second command returns a list of your network adapters that are RDMA capable. In this case, it should show your RoCE-capable cards only. The same command will work if you are using iWARP network adapters. The final command returns more detailed hardware information.

```
Get-NetOffloadGlobalSetting | Select NetworkDirect
Get-NetAdapterRDMA
Get-NetAdapterHardwareInfo
```

4. With the network adapter configuration verified, use the SMB Share cmdlets in Windows PowerShell to verify that SMB Multichannel is enabled, that the network adapters are properly recognized by SMB, and that the RDMA capability of the network adapters is properly identified. Run the following two commands, one after another. The first command check confirms that SMB Multichannel is enabled. With SMB Multichannel, SMB detects whether a network adapter has the RDMA capability and then creates multiple RDMA connections for that single session (two per interface). This allows SMB to use the high-throughput, low-latency, and low-CPU utilization that RDMA-capable network adapters offer. It also offers fault tolerance if you are using multiple RDMA interfaces. The second command returns more information about the network adapters.

```
Get-SmbClientConfiguration | Select EnableMultichannel
Get-SmbClientNetworkInterface
```

5. Repeat steps 1 through 4 of this procedure on MGMT02, HV01, HV02, HV03, HV04, FS01, and FS02.

6. Log on to FS01 using contoso\administrator credentials.
7. Open a Windows PowerShell window as an administrator.
8. Run the following three commands, one after another. Similar to before, the first command checks to ensure that SMB Multichannel is enabled, this time from the server side. The second command provides more detail about the network adapters within the server. The final NETSTAT command confirms whether the file server is listening on the RDMA interfaces.

```
Get-SmbServerConfiguration | Select EnableMultichannel
Get-SmbServerNetworkInterface
netstat.exe -xan | ? {$_ -match "445"}
```

9. Repeat steps 6 through 8 of this procedure on FS02.
10. To verify the SMB connection and that RoCE is operating as expected, log on to MGMT01 using contoso\administrator credentials. Open an administrative Windows PowerShell window.
11. Start a long-running file copy to create a lasting session with the SMB server, in this case, the SOFS. For this test, you could use your SYSPREP.vhdx file located on D:\Exports\Sysprep\Virtual Hard Disks. You could also use some of the ISO files within D:\Software. While the copy is ongoing, run the following cmdlets to verify the connection is using the correct SMB dialect and that SMB Direct is working. Note that if you have no activity while you run the commands above, it's possible you'll get an empty list. This would likely be because your session has expired, and there are no current connections. When running the third command, you should see multiple active connections of type Connection and Listener, all under Kernel mode.

```
Get-SmbConnection
Get-SmbMultichannelConnection
netstat.exe -xan | ? {$_ -match "445"}
```

12. Check performance within Performance Monitor, focusing specifically on the counters for SMB Direct Connection and RDMA Activity.
13. Additionally, in Task Manager, view the Performance tab and select one of your RoCE-capable network adapters. You should see minimal TCP/UDP send and receive speeds, even though the file copy is taking place. This is because the file copy is offloaded to the network adapter.
14. Get more information in the event log, with details on how to enable this on TechNet at http://technet.microsoft.com/en-us/library/ad08f159-9433-4a27-81ab-02b8c030939a#BKMK_verify_config.

Procedure 27: Test overall storage health (optional)

With the storage connectivity configured and optimized, one final, optional step is to take advantage of resources published by the Microsoft storage engineering team. Specifically these resources help assess the health of the overall storage infrastructure.

To download the Storage Cluster Health Test, visit the TechNet gallery at <https://gallery.technet.microsoft.com/scriptcenter/Test-StorageHealthps1-66d84fd4>. This script checks the health and capacity of a storage cluster based on Windows Server 2012 R2 SOFS. The script checks various parts of the cluster, including the cluster resources, networks, nodes, disks, enclosures, virtual disks, CSVs, file shares, and more. There is also a useful performance testing capability.

Procedure 28: Configure Live Migration over SMB

With RoCE deployed and tested, the final optimization step is to configure the Live Migration feature on each of your Hyper-V hosts to use the Live Migration over SMB option instead of the default Live Migration with compression. Live Migration over SMB enables you to perform a live migration of VMs by using SMB 3.0 as a transport. This allows you to take advantage of key SMB features such as SMB Direct and SMB Multichannel by providing high-speed migration with low CPU utilization.

1. Log on to VMM01 using contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, and connect to VMM-HA.
3. With VMM-HA opened, click the Fabric workspace, expand All Hosts, expand COMPUTE, and click HVCLUSTER.
4. Right-click HV01, and click Properties.
5. In the HV01 Properties window, click Migration Settings.
6. Under Performance Option, select the Use SMB As Transport option, as shown in Figure 5-26, and then click OK.




FIGURE 5-26 Selecting the Use SMB As Transport option

7. Repeat steps 4 through 6 of this procedure for HV02, HV03, and HV04.
8. Expand the MGMT host group, and then click MGMTCLUS.
9. Repeat steps 4 through 6 of this procedure for MGMT01 and MGMT02.

Changing this setting on each of the hosts means that live migrations of VMs to and from the hosts will use the SMB protocol. With RoCE configured in the previous steps, that traffic is offloaded to the network adapters, reducing the impact to the host while accelerating live migration traffic in contrast with the previous, with compression, option.

Configuring network virtualization

When you've completed the steps in Chapter 5, "Configuring compute infrastructure," you've successfully deployed your back-end Scale-Out File Server (SOFS) storage and your Hyper-V cluster and set up your management cluster to manage them. At this point, you are ready to begin deploying virtual machines (VMs) on top of your Hyper-V compute cluster.

However, as the deployment now stands, when you start to deploy new virtual workloads, you'll have only one usable VM network for VM communication. This network was automatically created in Chapter 3, "Configuring network infrastructure," when you created your underlying Tenant_LN logical network. Recall that when creating this Tenant_LN logical network, you allowed future VM networks created on top of this logical network to use network virtualization. You also selected the option to automatically create a VM network with the same name. This VM network was created to allow VMs attached to this VM network to communicate directly on the underlying logical network.

For example, imagine you had a physical server somewhere on your network with the IP address 10.10.0.253. Your Microsoft System Center Virtual Machine Manager server, which is inside a VM, has the IP address 10.10.0.13 and is attached to the VM network Tenant_LN. Allowing the Tenant_LN VM network to communicate directly through the Tenant_LN logical network enables your VMM01 server and the physical server to communicate successfully.

In the POC, for the management VMs that reside on the management cluster, this configuration is acceptable. However, suppose you want to deploy other workloads onto, for example, the Hyper-V cluster (HV01 through HV04). If you place those workloads on the Tenant_LN VM network also, they too would have access to the underlying logical network. From a security and isolation perspective, that access may not be desirable.

To address such situations, one method of isolation would be the use of virtual LANs (VLANs). You could create a new VM network and specify a VLAN as part of that creation process. Any VMs that are created and assigned to that VM network would be isolated from other VMs, including the management VMs because they are on a separate VLAN from the Tenant_LN VM network. Although VLANs are a simple concept, as discussed in Chapter 1, "Design and planning," they can be cumbersome to manage. This is certainly true at scale, which is where Hyper-V network virtualization can provide additional flexibility and simplicity and huge scalability.

With Hyper-V network virtualization, as discussed in Chapter 1, you can isolate individual VMs or groups of VMs from one another by assigning them to different VM networks. Hyper-V network virtualization isolates VM networks from one another but allows VMs on the same VM networks to successfully communicate with one another. The challenge comes, however, when VMs within these isolated VM networks need to communicate out of their VM networks—perhaps they need to connect with a physical server elsewhere on the network, or perhaps a developer’s workstation needs to reach the isolated virtualized test lab. The physical server or the developer’s workstation would not understand Hyper-V network virtualization, nor would they be able to reach those VMs within the VM networks unless you also deployed Windows Server Gateway. As discussed in more depth in Chapter 1, the gateway provides several different functions. The core function of the gateway in this POC is to bridge the non-virtual and virtual-networked environments. A common requirement that the gateway meets is to allow VMs within a VM network to communicate out to the Internet. Figure 6-1 shows the architecture of the Windows Server Gateway configuration.

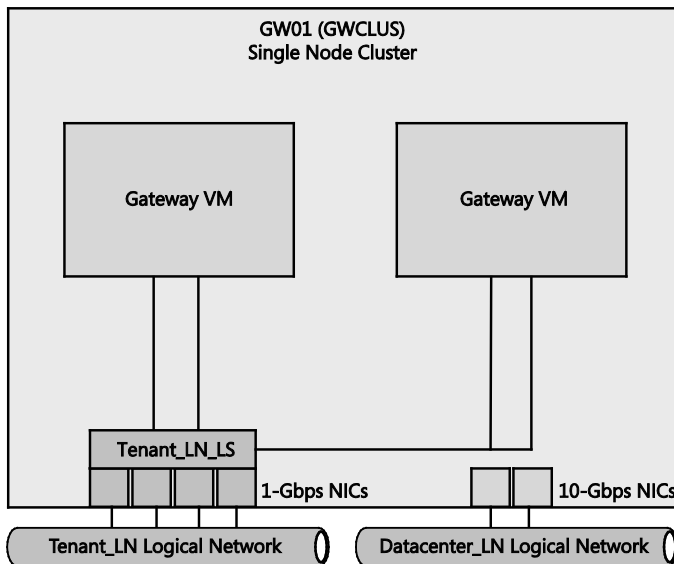


FIGURE 6-1 An architectural representation of the Windows Server Gateway configuration

In this chapter, you’ll start by configuring a new Hyper-V host. This Hyper-V host will be dedicated to running the Windows Server Gateway VMs. In a production environment, you would deploy multiple physical hosts configured as a cluster to act as dedicated Windows Server Gateway Hyper-V hosts. For this POC configuration, a single-host cluster will be fine.

To deploy this host, you’ll build on skills you learned earlier, specifically the ability to perform a bare-metal deployment. After the Hyper-V host is deployed, you’ll make some minor configuration changes to the new host to ensure it’s ready to act as the new Windows Server Gateway Hyper-V host.

When the Windows Server Gateway host is configured, as shown in Figure 6-1, you'll create and deploy several new VMs that will become the Windows Server Gateway VMs. These VMs, managed by System Center Virtual Machine Manager, interact with the host and enable VMs deployed with network virtualization to communicate correctly.

With the Windows Server Gateway VMs deployed, configured, and under System Center Virtual Machine Manager management, you'll use network virtualization to create and deploy several new VM networks. Into these VM networks, you will deploy several VMs that you will use to test communication within and across VM networks.

By the end of this chapter, you will have covered all of the key aspects of deploying Hyper-V with software-defined storage and networking and will be able to move on to testing some of your own workloads within the configuration.

Configuration walkthrough

In the following sections, you walk through the key procedures and steps outlined in the previous section. You begin by preparing the underlying physical infrastructure in advance or by using System Center Virtual Machine Manager to deploy a new bare-metal Hyper-V host that will ultimately become your Windows Server Gateway Hyper-V host.

In this configuration, the Hyper-V host will be a Dell PowerEdge R620 server with the following specifications. Your server's specifications may vary.

- 128-GB RAM, dual-socket, 2-Ghz Intel Xeon E5-2650 processor, each with eight cores
- Two 10-Gbps RoCE-capable NICs, used for DatacenterNetwork 1 and DatacenterNetwork 2
- Four 1-Gbps NICs, used for TenantNetwork to allow the Windows Server Gateway VMs to communicate over the physical network and to bridge the virtual and non-virtual network traffic
- One BMC port for lights-out management over the network

In general, a dual-socket server-class platform provided by a major Original Equipment Manufacturer (OEM) should be a sufficient base platform. The amount of RAM in the server typically determines the number of VMs that can effectively run on the server appliance.

The choice of network interfaces can have a great impact on the performance and scalability of the gateway appliance. Network interface offloads substantially benefit the Hyper-V switch performance. A new feature in Windows Server 2012 and Windows Server 2012 R2 is Generic Route Encapsulation Task Offload. With it, the network interface will operate on the inner packet header for standard offloads so that offloads such as Large Send Offload (LSO) and Virtual Machine Queue (VMQ) benefit performance and scalability. Bear these key characteristics in mind when you're procuring hardware that will operate within a network virtualization environment.

Procedure 1: Rack and connect the Windows Server Gateway Hyper-V host

In this procedure, you physically deploy the hardware into the environment. This may involve the acquisition of new hardware or repurposing of existing hardware. If you already have hardware racked and cabled in your environment, ensure that the networking is configured according to the instructions in this procedure. Name the host GW01.

1. Rack the physical server that will become GW01.
2. Connect power. Connect the network, according to the following directions. Note that at this point, you are just plugging in the hardware. The network configuration will be automatically completed at deployment time.

The Windows Server Gateway Hyper-V host will use three separate networks:

- DatacenterNetwork Port 1 (10 Gbps, not teamed)
- DatacenterNetwork Port 2 (10 Gbps, not teamed)
- TenantNetwork Team (teamed, consisting of four 1-Gbps adapters)

Plug all NICs on each node into physical network ports across separate switches:

- DatacenterNetwork Port 1 on each node to a switch that is RoCE-capable
- DatacenterNetwork Port 2 on each node to another switch, also RoCE-capable
- Half of the ports for TenantNetwork team to one switch and half to another switch for redundancy

Procedure 2: Configure BMC

In this chapter, as mentioned earlier, you'll use System Center Virtual Machine Manager to deploy the Windows Server Gateway Hyper-V host and take advantage of the bare-metal provisioning capabilities that you used to deploy the SOFS and Hyper-V clusters earlier. A key enabler of this bare-metal deployment capability is the use of the baseboard management controller (BMC). The BMC enables out-of-band management and allows System Center Virtual Machine Manager to power on the physical server, discover it, control the network boot, and ultimately power off the physical server—all before the installed operating system starts. For this to work, you need to have the BMC within each host configured in advance to ensure System Center Virtual Machine Manager can communicate with them.

Table 6-1 shows the BMC settings for the Windows Server Gateway Hyper-V host. Note that each hardware vendor's BMC configuration utility may differ from that of other vendors, so please refer to your chosen hardware vendor's guidance on how to configure your BMC. After the BMC is configured, note the username, password, and IP address.

TABLE 6-1 BMC settings for Windows Server Gateway Hyper-V host

PHYSICAL NIC	PURPOSE	ADDRESSES
GW01 – BMC Network Adapter	BMC interface for PXE Boot	IP: 10.10.0.7 SM: 255.255.255.0 DG: 10.10.0.254

For a successful bare-metal deployment, your Windows Server Gateway Hyper-V host must support one of the following out-of-band management protocols:

- Intelligent Platform Management Interface (IPMI) versions 1.5 or 2.0
- Data Center Management Interface (DCMI) version 1.0
- System Management Architecture for Server Hardware (SMASH) version 1.0 over WS-Management (WS-Man)

Whichever protocol you are using, ensure that you have the latest version of firmware for your BMC model.

Procedure 3: Configuring a BMC administrator in System Center Virtual Machine Manager (Optional)

As mentioned earlier, for System Center Virtual Machine Manager to successfully control the Windows Server Gateway Hyper-V host through the BMC, you will need to give System Center Virtual Machine Manager the specific credentials for the BMC for the purpose of bare-metal deployment. You can easily achieve this in System Center Virtual Machine Manager by specifying a new run-as account that contains the credentials required for the BMC.

Note that if you are using the same credentials for the Windows Server Gateway Hyper-V host BMC as you used for the SOFS and Hyper-V cluster BMCs, skip this step, since you already have an appropriate run-as account configured. If you are using different credentials, proceed with this step.

1. Log on to your VMM01 VM using your contoso\administrator credentials.
2. From the desktop, launch the System Center Virtual Machine Manager console. For the name, enter **VMM-HA**, and click Connect. By entering VMM-HA, you'll be connecting to the highly available System Center Virtual Machine Manager configuration you constructed in Chapter 2, "Deploying the management cluster."
3. In the bottom-left corner of the System Center Virtual Machine Manager console, click Settings.
4. Expand Security, and select Run As Accounts. On the top ribbon navigation, select Create Run As Account.
5. In the Create Run As Account wizard, in the Name text box, enter **GW BMC Administrator**. Optionally, add a description.
6. In the User Name text box, enter the user name that is relevant for your BMC settings.

7. Below the user name, enter a password and password confirmation in the corresponding text boxes.
8. Unless your BMC account is within your domain, clear the Validate Domain Credentials check box and click OK.

You now have a run-as account that corresponds to the BMC on your gateway node.

Procedure 4: Create a physical computer profile

As you learned in Chapter 4, “Configuring storage infrastructure,” and in Chapter 5, physical computer profiles define the standardized characteristics of a physical server deployment that results in deployment of either a Hyper-V host or an SOFS. The concept of using a physical computer profile for physical server deployment is similar to the concept of using VM templates for VM deployment. The following procedure describes how to create a physical computer profile, this time in the System Center Virtual Machine Manager library and specifically for the Windows Server Gateway Hyper-V host rather than for the regular Hyper-V or SOFS cluster nodes, which you saw in previous chapters.

1. In the bottom-left corner of the System Center Virtual Machine Manager console, click Library.
2. Right-click Physical Computer Profile and select Create Physical Computer Profile, as shown in Figure 6-2.




FIGURE 6-2 Context menu for creating a physical computer profile

3. On the Profile Description page, name the profile **Network Virtualization Gateway Hyper-V Host**. Optionally, set a description. Ensure that the VM Host option is selected, as shown in Figure 6-3, and then click Next

A screenshot of the 'Provide a name for the physical computer profile' dialog box. The 'Name' field contains 'Network Virtualization Gateway Hyper-V Host'. The 'Description' field is empty. Under the 'Role' section, the 'VM Host' radio button is selected, and the 'Windows File Server' radio button is unselected.

FIGURE 6-3 Naming the profile

4. On the OS Image page, shown in Figure 6-4, click Browse and select the VHDX file that you published to the System Center Virtual Machine Manager library in Chapter 2. Optionally, you can select the Do Not Convert The Virtual Hard Disk Type To Fixed During Deployment check box. For this POC configuration, select the option, and then click Next.




FIGURE 6-4 Specifying settings on the OS Image page

5. On the Hardware Configuration page, under Management NIC, click IP Configuration.
6. Select Allocate A Static IP From The Following Network, and select the DataCenter_LN logical network from the drop-down list, as shown in Figure 6-5.




FIGURE 6-5 Allocating a static IP address

7. At the top of the window, click Add, and then select Physical Network Adapter.
8. Repeat steps 5 through 6 of this procedure for Physical NIC #1.
9. At the top of the window, click Add, and then select Physical Network Adapter.
10. Under Physical NIC #2, click Connectivity Properties.
11. Select the Connect This Physical NIC To The Following Logical Switch check box. From the drop-down list, select the logical switch you created in Chapter 3. This will automatically select the uplink port profile that was created in Chapter 3.
12. Repeat steps 9 through 11 of this procedure, adding another three physical NICs. When this is complete, you will have a Management NIC and Physical NIC #1 through #5. Two of the NICs will represent the 10-Gbps DatacenterNetwork adapters, and the remaining four NICs will be transformed into a logical switch, allowing Windows Server Gateway VM traffic to flow onto the physical network.
13. Under Disk And Partitions, click OS, and review the settings. You will be using 100 percent of the available local disk for the deployment, but you can adjust if required.
14. Under Driver Options, click Driver Filter. System Center Virtual Machine Manager provides the ability to include hardware-specific drivers as part of the deployment process. This section of the wizard assumes you have already imported the drivers into the System Center Virtual Machine Manager library. This will not be covered as part of this configuration; you will use the built-in Windows Server drivers for the time being. Click Next.

15. On the OS Configuration page, perform the following steps:
 - Type **contoso.com** as the domain to join, and select the SetupAdmin run-as account as the credentials for joining the domain.
 - For the Admin Password, enter the standard password you're using for the POC. This will be used to set the local administrator account password on the server.
 - For Identity Information, enter your company information (optional).
 - For Product Key, enter your product key if you are using one for the POC (optional).
 - Select your time zone.
 - Select an answer file if you're using one for any of your Windows settings (optional).
16. Click Next. On the Host Settings page, leave the paths blank for now; you will attach the file shares when the hosts have been deployed. Click Finish.
17. Monitor the job for successful completion, and then close the Jobs window.

Procedure 5: Create a gateway host group

Before you discover and provision your new Windows Server Gateway Hyper-V host, create a new container for this host within System Center Virtual Machine Manager. This will aid with organization. As this will be a single host within the host group rather than a cluster, policies you ordinarily define at the host group level, such as Dynamic and Power Optimization, will not apply.

1. Log on to your VMM01 VM using your contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, and then click Fabric.
3. Expand Servers, and then right-click All Hosts and select Create Host Group.
4. Enter the name **GATEWAY** and click outside the text box. Your new host group is created.

Procedure 6: Discover and provision the Windows Server Gateway Hyper-V host with System Center Virtual Machine Manager

You've defined the physical computer profile, the VHDX is ready in the System Center Virtual Machine Manager library, and you configured the Windows Deployment Service (WDS) server in Chapter 2. Now you're can use System Center Virtual Machine Manager to deploy your Windows Server Gateway Hyper-V host.

As a recap, when you create new Hyper-V hosts from bare-metal computers, at a high level, the Add Resources Wizard does the following:

- Discovers the physical computer through out-of-band management

- Deploys an operating system image on the computer through the host profile or the physical computer profile
- Enables the Hyper-V role on the computers
- Brings the computer under System Center Virtual Machine Manager management as a managed Hyper-V host

Upon completion of the wizard, you need to perform additional steps to enable this new Hyper-V host to function correctly as a Windows Server Gateway Hyper-V host because these steps are not covered within the bare-metal deployment process. The following steps walk through the deployment of Hyper-V to your bare-metal machine.

1. Log on to your VMM01 VM using your contoso\administrator credentials.
2. Open the System Center Virtual Machine Manager console, and then click Fabric.
3. In the Fabric pane, click Servers.
4. On the Home tab, in the Add group, click Add Resources, and then click Hyper-V Hosts And Clusters. The Add Resources Wizard opens.
5. On the Resource location page, click Physical Computers To Be Provisioned As Virtual Machine Hosts, and then click Next.
6. On the Credentials And Protocol page, next to the Run As account box, click Browse, select your BMC administrator account, and then click OK.
7. In the Protocol list, click IPMI On Port 623 For Discovery, and then click Next. This selection may vary depending on your hardware.
8. On the Discovery Scope page, depending on how your BMCs were configured, enter the IP address directly or select either the IP subnet or, more specifically, IP range. Whichever you choose, System Center Virtual Machine Manager will scan the infrastructure and return a list of discovered hosts. Hosts do not have to be powered on to be discovered. Click Next to start discovery.
9. Once discovery is completed, on the Target Resources page, select the discovered host, and click Next.
10. On the Provisioning Options page, in the Host Group list, select GATEWAY as the target location for the new Windows Server Gateway Hyper-V host, and from the drop-down list select your Network Virtualization Gateway Hyper-V host physical computer profile. Click Next.
11. On the Deployment Customization page, notice that deep discovery is running. System Center Virtual Machine Manager will have used the BMC to wake the server and will begin to capture details about its configuration, components, and more. This will take a few minutes. When this is completed, the server returns to a powered-off state. With the process complete, address the warnings that appear before you proceed. Select the server by selecting the BMC IP address with the warning triangle.
12. In the Computer Name box, enter **GW01**, as shown in Figure 6-6, and then click Network Adapters.

BMC IP address: 10.10.0.7
 SMBIOS ID: 4c4c4544-0039-3610-8038-c3c04f565231
 Serial number: C968VR1
 Manufacturer: DELL
 Computer name: GW01
 Skip Active Directory check for this computer name

FIGURE 6-6 Entering the computer name

13. You should see six network adapters listed, as in Figure 6-7. If the window is cramped with information, stretch it to the right by clicking the edge of the window and dragging until you are happy with the size.

In this configuration, two of the network adapters are 10 Gbps and will be placed onto the DatacenterNetwork. The remaining four will be aggregated into a team and have a logical switch deployed, allowing the to-be-deployed Windows Server Gateway VMs to communicate out onto the network.

MAC Address	Name	Logical Swit...	IP Assignm...	Management NIC
00:0A:F7:5D:59:A0	SLOT 2 Port 1	None	Static IP	Yes
00:0A:F7:5D:59:A2	SLOT 2 Port 2	None	Static IP	No
E0:DB:55:20:A1:74	NIC1	Tenant_LN	N/A	No
E0:DB:55:20:A1:75	NIC2	Tenant_LN	N/A	No
E0:DB:55:20:A1:76	NIC3	Tenant_LN	N/A	No
E0:DB:55:20:A1:77	NIC4	Tenant_LN	N/A	No

FIGURE 6-7 Six network adapters listed

14. Ensure that one of your 10-Gbps network adapters under Management NIC has Yes selected. Also ensure that both of the 10-Gbps adapters have Static IP selected under IP Assignment. If you are unsure which of the adapters your 10-Gbps adapters are, note the MAC address and refer to your documentation, or use a web browser to visit the BMC management IP address and use the BMC configuration page to confirm the network adapter details.
15. For the first of your 10-Gbps adapters, click the ellipses (...) button to further customize the NIC settings. This will open the Network Adapter IP Configuration window.
16. In the Network Adapter IP Configuration window, select the Specify Static IP Settings check box for this network adapter.
17. For logical network, use the drop-down list to select Datacenter_LN.
18. For IP subnet, use the drop-down list to select 10.10.1.0/24, and select the Obtain An IP Address check box corresponding to the selected subnet. As part of the deployment process, System Center Virtual Machine Manager will assign a static IP address from the pool you created in Chapter 3. Click OK.
19. Repeat steps 15 through 18 of this procedure, but this time select the 10.10.2.0/24 subnet.

20. The four remaining 1-Gbps adapters in this configuration will be used exclusively by VMs. Therefore, as part of the deployment, specify that these network adapters should have a logical switch assigned to them. For each of the 1-Gbps network adapters, in the Logical Switch column, ensure they all have the Tenant_LN_LS selected.
21. For the first of your 1-Gbps adapters, click the ellipses (...) button. This will open the Network Adapter IP Configuration window.
22. In the Network Adapter IP Configuration window, select the Connect This Physical NIC To The Following Logical Switch check box: Using the drop-down list, ensure that Tenant_LN_LS is selected. This is the logical switch that was constructed in Chapter 3.
23. Under Apply The Following Uplink Port Profile To This Physical NIC, use the drop-down list to select the uplink port profile. You created only a single profile in Chapter 3, so this will be the only one available for selection, as shown in Figure 6-8. Click OK.

The screenshot shows a configuration window with two sections. The first section is titled "Connect this physical NIC to the following logical switch:" and has a checked checkbox. Below it is a dropdown menu with "Tenant_LN_LS" selected. The second section is titled "Apply the following uplink port profile to this physical NIC:" and has a dropdown menu with "Tenant_LN_UPP_c830e941-cb8f-4407-b905-a4f83d7c9c11" selected.

FIGURE 6-8 Selecting the uplink port profile

24. Repeat steps 21 through 23 of this procedure for the remaining 1-Gbps network adapters.
25. Click Disks. If you have more than one local disk in the target server, you can use the check box and drop-down list to select the disk you'd like to deploy Hyper-V onto.
26. On the Summary page, review your customization settings, and then click Finish.

System Center Virtual Machine Manager will then begin the deployment process, transferring a copy of the virtual hard disk, which is currently stored in the System Center Virtual Machine Manager library, to the Windows Server Gateway Hyper-V host. This will configure the host to natively boot from the virtual hard disk. Then System Center Virtual Machine Manager will proceed with the automated setup of the Windows Server Gateway Hyper-V host, including joining the domain, enabling Hyper-V, and deploying the System Center Virtual Machine Manager management agent. When this is completed, you will see the new Windows Server Gateway Hyper-V host under System Center Virtual Machine Manager management within the GATEWAY host group.

Procedure 7: Update drivers and firmware on GW01

With your new Hyper-V host deployed, it's important to check the firmware for the key hardware components of the host and update them if necessary. If you have acquired new hardware, it is less likely this step will be necessary. But if you are repurposing existing hardware, this step ensures that performance and reliability of the configuration is high. Your main focus should be on the BIOS and the BMC, along with the firmware and drivers for your

network adapters—both 1-Gbps and the 10-Gbps adapters—on the new Windows Server Gateway Hyper-V host.

Refer to your hardware vendor's guidance for obtaining and applying any firmware updates. Once you've applied these and rebooted the server, you can proceed with the next procedure. Also, if it has been some time since you created your base SYSPREP.vhdx file and deployed your Windows Server Gateway host, check for the latest Windows updates and reboot as necessary.

It's important to note that in a production environment, you'd typically deploy multiple, dedicated Hyper-V hosts to provide the required Windows Server Gateway capabilities, but for the purpose of this POC configuration, a single, dedicated Hyper-V host is fine.

Procedure 8: Configure GW01 to run Windows Server Gateway VMs

With your new Windows Server Gateway Hyper-V host assigned with resilient, high-performing shared storage, you need to apply a few more minor changes to finalize the host so it is ready to run the Windows Server Gateway VMs. The first change you need to make is to add the host to a cluster. When you deploy the Windows Server Gateway VMs in a later procedure, you will be using a pre-constructed service template provided by Microsoft. This service template, as part of the final deployment steps, requires an underlying Hyper-V host cluster. In a production environment, the Windows Server Gateway VMs should be deployed to a resilient Hyper-V cluster, consisting of multiple physical hosts. In this configuration, however, you have only GW01, a single Hyper-V host. As a result, you will still create the cluster, but it will be a single-node cluster.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Fabric workspace.
2. Expand Servers, then All Hosts, and select GATEWAY.
3. On the ribbon, click Create, and then select Hyper-V Cluster. The Create Cluster Wizard opens.
4. On the General page, for Cluster Name, enter **GWCLUS**.
5. Ensure that the Use An Existing Run As Account option is selected, and click Browse. Select SetupAdmin, click OK, and then click Next.
6. On the Nodes page, use the Host Group drop-down list to select GATEWAY. Under Available Hosts, select gw01.contoso.com, click Add, and then click Next.
7. On the IP Address page, ensure both check boxes under Network are selected, and under Static IP Pool, use the drop-down list for each of the networks to select DatacenterNetwork Pool1 and DatacenterNetwork Pool2 respectively. Click Next.
8. On the Storage page, click Next.
9. On the Virtual Switches page, click Next.
10. On the Summary page, review your settings and click Finish.

11. Monitor the job for successful completion, and then close the Jobs window.
12. Return to the Fabric workspace, expand Servers, then All Hosts, and then GATEWAY. Right-click GWCLUS and click Properties.
13. On the General tab, next to Cluster Reserve (Nodes), enter **0**.
14. On the File Share Storage tab, click Add. In the Add File Share window, shown in Figure 6-9, use the drop-down list to select InfraShare1, and then click OK. Because the VMs that will run on this Windows Server Gateway Hyper-V host will be providing infrastructure services, it makes sense to deploy them onto InfraShare1 rather than one of the tenant shares.




FIGURE 6-9 Selecting InfraShare1

15. Repeat step 14 of this procedure to add the InfraShare2, and then click OK to close the GWCLUS Properties window.
16. With GWCLUS selected, in the central pane, right-click GW01, and select Properties. Click the Hardware tab.
17. Under Network Adapters, scroll down until you find your first 10-Gbps network adapter. Click the network adapter name. In the details panel to the right, you should see more information, including the IPv4 address. It should be either 10.10.1.x or 10.10.2.x.
18. Below the name of the network adapter, click Logical Network. Select the check box that applies to that network adapter. For instance, if the network adapter you selected had an IPv4 address of 10.10.1.x, select the Datacenter_LN first subnet, 10.10.1.0/24.
19. Repeat steps 17 through 18 of this procedure for your remaining 10-Gbps network adapter, this time selecting the other Datacenter_LN logical network subnet.
20. Click the Host Access tab. Select the This Host Is A Dedicated Network Virtualization Gateway check box, shown in Figure 6-10. As a result of choosing this option, the host is not available for placement of VMs that require network virtualization.




FIGURE 6-10 Configuring GW01 to be a dedicated network virtualization gateway host

21. With your settings completed, click OK to close the GW01 Properties window.

Procedure 9: Rename existing VM network

In Chapter 3, you defined two logical networks. These logical networks were created to represent the underlying physical networks that would be used within the infrastructure. The Datacenter_LN logical network was created to represent the physical networks that would handle storage, live migration, cluster, and management traffic and would not support VM traffic. For VM traffic, you created the Tenant_LN logical network. This Tenant_LN logical network would represent the underlying physical network on which traffic in and out of VMs would travel. When you created this logical network, you also created a VM network with the same name to allow VMs to access this logical network directly. System Center Virtual Machine Manager subsequently created a corresponding Tenant_LN VM network that can be used by VMs. You also created an IP pool for the underlying Tenant_LN logical network. This pool had a range from 10.10.0.100 through to 10.10.0.250.

With System Center Virtual Machine Manager now managing your management hosts, the VMs that are running on those hosts will already have been assigned to the Tenant_LN VM network. However, for simplification and ease of identification, in this step you rename the Tenant_LN VM network to something that better represents the use of the VM network at this point.

1. On VMM01, in the System Center Virtual Machine Manager console, open the VMs And Services workspace.
2. Click VM Networks. Only one VM Network will appear in the central pane.
3. Right-click Tenant_LN and select Properties.
4. In the Tenant_LN Properties window, next to Name, enter **Management VM Network**, and click OK.

As it stands, the only VMs that are technically attached to this VM network are the management and infrastructure VMs. The hosts themselves, across management, compute, and storage, all reside on the Datacenter_LN logical network and do not require a VM network. The VMs that reside on MGMTCLUS—such as VMM01, VMM02, SQL01, and so on—provide management functions and should therefore reside on an appropriately named VM network.

When you deploy new tenant VMs later in this chapter, they will be deployed to their own newly created VM networks, not the Management VM Network that you've just modified.

Procedure 10: Download and import the Windows Server Gateway service templates into System Center Virtual Machine Manager

With the host configured, you're now in a position to construct your Windows Server Gateway VMs. These VMs are Windows Server 2012 R2 VMs configured with the Routing and Remote Access (RRAS) role and accompanying features. With the correct configuration, the Windows Server instance is transformed into a multi-tenant software gateway and Border Gateway Protocol (BGP) router that allows cloud service providers and enterprises to enable datacenter and cloud network traffic routing between virtual and physical networks, including the Internet.

Rather than just deploy a standard VM, you will use System Center Virtual Machine Manager to deploy a service template. In System Center Virtual Machine Manager, a service is a set of VMs that are configured and deployed together and managed as a single entity—for example, a deployment of a multi-tier line-of-business application. The gateway you will be configuring is also a service that can be backed by one or, ideally, many VMs.

To assist with the deployment of the Windows Server Gateway VMs, Microsoft has made available a set of downloadable resources that you can quickly customize for your specific environment. In this procedure, you download the relevant files and place them in the System Center Virtual Machine Manager library, ready for use in the next procedure.

Within the download files, you will notice two service templates, denoted by the .xml file extension. These service templates represent two different deployment types. The first is for a Windows Server Gateway VM with two virtual network adapters, and the second is for a Windows Server Gateway VM with three virtual network adapters. Typically, a Windows Server Gateway that is used for network virtualization requires three types of network:

- **A back-end network for network virtualization** This is the network that Hyper-V network virtualization uses to send encapsulated packets to and from the tenant VMs. To maintain isolation, this network must only be used for Hyper-V network virtualization.
- **A front-end network** This is the network that serves as the external side of the network through which your virtual networks can access outside physical networks. This network must have an Internet-routable IP address space if you're using site-to-site VPN functionality. System Center Virtual Machine Manager must have a static IP address pool for this network so it can assign IP addresses for network address translation (NAT).
- **A management network** This is the network on which System Center Virtual Machine Manager communicates with the Hyper-V hosts and the Windows Server Gateway VMs. It must have a domain controller available and DNS registration for the guest cluster. You can use either static IP address assignment or DHCP on this network.

In this configuration, the back-end network is the Tenant_LN logical network, which you created in Chapter 3. Encapsulated traffic flows over this network between tenant VMs located on the Hyper-V compute cluster and through the Windows Server Gateway VMs. The front-end network, required to reach outside networks including the Internet and also the management network in this configuration are the same network. VMs with a 10.10.0.x IP address, such as the management VMs, can reach the Internet through the previously configured default gateway, configured with a 10.10.0.254 IP address. In this configuration, you will be using the service template for two network adapters.

For further information on the two- and three-virtual network adapter service template configuration, refer to the guidance within the Windows Server Gateway download files, which

you will download in step 1 below. You'll also need two empty disk images in your library for the cluster disks, which you create using Windows PowerShell.

1. On VMM01, download the compressed file (with a .zip extension) for the Windows Server Gateway from the Microsoft website at <http://go.microsoft.com/fwlink/p/?LinkId=329037>.
2. When the file is downloaded, extract the files in the download. These files include a Quick Start Guide, two service templates, and a custom resource folder (a folder with a .cr extension) that contains files required for the service templates.
3. On VMM01, in the System Center Virtual Machine Manager console, open the Library workspace.
4. Expand Library Servers, then library.contoso.com, right-click LibraryShare, and click Explore.
5. In the LibraryShare window, right-click under VHDs and select New, then Folder. Name the new folder **Windows Server Gateway**. Minimize File Explorer.
6. Log on to MGMT01. Right-click the Windows PowerShell icon on the taskbar, right-click Windows PowerShell, and select Run As Administrator.
7. Run the following commands:

```
New-VHD -Path "D:\Software\VM Gateway CSV.vhdx" -SizeBytes 10GB
New-VHD -Path "D:\Software\VM Gateway Quorum.vhdx" -SizeBytes 1GB
```
8. Return to VMM01, in the Library workspace, on the ribbon, click Import Physical Resource. The Import Library Resources window opens.
9. Click Add Custom Resource. Browse to the location of your extracted service template files, select the VMClusterSetup.cr folder, and click OK.
10. Click Add Resource. In the Select Resource Items window, in the address bar, navigate to MGMT01\Software. Select both of the VM Gateway virtual hard disk files, and then click Open.
11. Return to the Import Library Resources window and, under Select Library Server And Destination For The Imported Resources, click Browse.
12. Under library.contoso.com, select the Windows Server Gateway folder, click OK, and then click Import to import the files into the System Center Virtual Machine Manager library.
13. Return to your LibraryShare File Explorer window. All three files imported should be listed, as shown in Figure 6-11.




FIGURE 6-11 Importing template-related files into the new Windows Server Gateway folder

14. In the upper-left corner, expand Templates, and select Service Templates.
15. On the ribbon, select Import Template.
16. In the Import Package Wizard, on the Select Package page, click Browse.
17. Navigate to the directory where you extracted the Windows Server Gateway files from the download. Select the template file entitled Windows Server 2012 R2 HA Gateway 2NIC, click Open, and then click Next.
18. On the Configure References page shown in Figure 6-12, there are a number of items under Type: Library Resources. These are currently defined in the imported template and need to be mapped to objects within your library. Because they are currently not mapped, they show a warning triangle for each of them. To simplify the view, click and drag the bottom-right of the window.

Resource Name	Release	Usage	Current Mapping	
Hyper-V		Windows Server 2012 R2 Gateway	Hyper-V	
Type: Library Resources				
VMClusterSetup.cr		Windows Server 2012 R2 Gateway	VMClusterSetup.cr	
Windows 2012 R2 VHD or VHDX		Windows Server 2012 R2 Gateway	WS2012R2	
CSV.vhdx		Windows Server 2012 R2 Gateway	VM Gateway CSV.vhdx	
Quorum.vhdx		Windows Server 2012 R2 Gateway	VM Gateway Quorum.v...	

FIGURE 6-12 Mapping completed within the Import Package Wizard

19. Click the pencil icon at the end of the VMClusterSetup.cr row of the table. In the Import Package Wizard window, select your VMClusterSetup.cr that is located within your library, and click OK. You will see the warning triangle disappear.
20. Click the pencil icon at the end of the Windows 2012 R2 VHD or VHDX row of the table. In the Import Package Wizard window, select the WS2012R2 VHDX file that is located within your library, and click OK. The warning triangle disappears.
21. Click the pencil icon at the end of the CSV.vhdx row of the table. In the Import Package Wizard window, select the VM Gateway CSV.vhdx file that is located within your library, and click OK. The warning triangle disappears.
22. Click the pencil icon at the end of the Quorum.vhdx row of the table. In the Import Package Wizard window, select the VM Gateway Quorum.vhdx file that is located within your library, and click OK. The warning triangle disappears. Click Next
23. On the Summary page, review the settings and click Import. The Jobs window appears and shows the job status. Make sure that the job has a status of Completed, and then close the dialog box.

Procedure 11: Customize the Windows Server Gateway service template

With the Windows Server Gateway service template now imported, you need to apply a few final customizations to ensure this imported service template is configured optimally for your environment.

1. On VMM01, in the System Center Virtual Machine Manager console, open the Library workspace.
2. Expand Templates, and then select Service Templates. A single service template is listed. Select this service template.
3. On the ribbon, click Open Designer to open the Service Template Designer. This will allow further, more granular customization of the service template.
4. The Service Template Designer opens, and the Windows Server 2012 R2 Gateway 2NIC service template is represented in the central canvas.
5. Right-click the box with the blue VM icon and the title Windows Server 2012 R2 Gateway, and then select Properties. The Windows Server 2012 R2 Gateway Properties window opens.
6. On the General tab, clear the Create An Availability Set For The Tier check box.

Recall that in Chapter 5 you created several availability sets for your System Center Virtual Machine Manager and SQL Server cluster VMs. An availability set is a System Center Virtual Machine Manager feature that ensures that related VMs are kept apart on clustered Hyper-V hosts. This ensures a whole service isn't taken down as a result of a single Hyper-V host suffering an outage. With this Windows Server Gateway configuration, you will be deploying onto a single-node Hyper-V cluster. As a result, the Windows Server Gateway VMs cannot possibly be kept apart on different nodes.

7. On the General tab, under the This Machine Can Be Scaled Out check box, reduce the default instance count to 1.

In this configuration, you will be deploying the Windows Server Gateway VMs to a single-node Hyper-V cluster rather than a multi-node Hyper-V cluster, which is recommended in a production environment. In that respect, it is not advisable to deploy multiple Windows Server Gateway VMs that make up the same service on the same Hyper-V host. You can safely deploy multiple Windows Server Gateway VMs to the same host, provided they are part of different deployed services.

8. Click the Hardware Configuration tab. Under General, select Memory. Dynamic is currently selected. Change the Maximum memory to 4,096 MB.
9. Click Bus Configuration, click IDE Devices, and select your virtual hard disk. To the right, click the Classification drop-down list, and select Infrastructure Storage.
10. Expand the SCSI Adapter 0. Two virtual hard disks should appear. Ensure both of these also have the Classification set to Infrastructure Storage. This ensures the virtual hard disks are optimally placed on the SOFS, rather than local storage, at deployment time.

11. Click the OS Configuration tab. Under General Settings, select Product Key. If you want to use an evaluation mode, clear the text box. The word None should appear under Product Key. If you plan to use a product key, do not delete the data in the text box. Click OK to close the Windows Server 2012 R2 Gateway Properties window.
12. Return to the Service Template Designer, and select the NIC 1 box. In the Properties area in the bottom window, next to IPv4 Address Type, use the drop-down list to select Static.
13. On the ribbon, select Save And Validate. Close the Service Template Designer.

Procedure 12: Configure deployment of the Windows Server Gateway service

With the service template now customized, you are ready to deploy the Windows Server Gateway service. The service template that was imported and subsequently modified contains several variables that you will be prompted to supply when you configure the deployment. Supplying these variables during deployment configuration rather than having them hard-coded into the service template configuration gives you additional flexibility. After you specify those variables, System Center Virtual Machine Manager handles the automated deployment of the VMs and the resultant Windows Server Gateway service.

In the previous procedure, you adjusted the default instance count of the service template from two to one. This value was initially set to two so that during deployment System Center Virtual Machine Manager would create two VMs for this Windows Server Gateway service, deploy them to different hosts on an underlying Hyper-V host cluster, and then, inside the VMs, create a guest cluster to provide redundancy and active/passive failover. In this deployment, with the customizations you have made, System Center Virtual Machine Manager will deploy a single VM to provide the relevant Windows Server Gateway services. The VM will still be deployed in a clustered configuration; however, it will be a single-node cluster inside the VM. System Center Virtual Machine Manager takes care of all of the configuration to ensure your single-node guest cluster is configured correctly and is ready to provide the important relevant Windows Server Gateway services upon completion.

1. On VMM01, navigate to the Fabric workspace.
2. Expand Network, and then click Logical Networks. In the central pane, right-click your Tenant_LN_Pool and click Properties.
3. In the Tenant_LN_Pool Properties window, click the IP Address Range tab.
4. In the IP Addresses To Be Reserved For Other Uses text box, enter **10.10.0.100**, **10.10.0.101** and click OK. These IP addresses will be used in the deployment of the Windows Server Gateway services.
5. Return to the System Center Virtual Machine Manager console, and navigate to the Library workspace. Expand Templates, and then click Service Templates. Select your single service template in the central pane.

6. On the ribbon, click Configure Deployment. The Select Name And Destination window opens, as shown in Figure 6-13.
7. In the Name text box, enter **Windows Server Gateway 001**.
8. From the Destination drop-down list, select the GATEWAY Host Group.
9. Under Network Settings, click the ellipsis (...) in the Value column.
10. In the Select A VM Network window, select Management VM Network, and click OK.

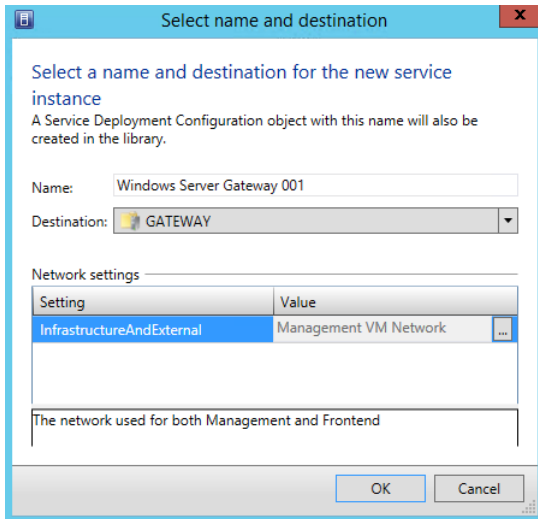


FIGURE 6-13 Configuring the deployment for the Windows Server Gateway service

11. Return to the Select Name And Destination window, and click OK to apply the settings.
12. The Deploy Service window opens. Maximize the window to facilitate data entry. It is normal for the VM instances in the central pane to initially have red or yellow icons. On the ribbon, click Refresh Preview to have the deployment service automatically find suitable hosts for the VMs.
13. In the bottom-left of the Deploy Service window are several other variables that you need to supply before the service can be successfully deployed, as shown in Figure 6-14. Where data is missing in the Value column and there is a red X, click the empty box and enter the information as indicated in the following list:
 - **DomainFQDN** Enter **contoso.com**.
 - **DomainUserRAA** Select the SetupAdmin run as account
 - **DomainUserRAAName** Enter **contoso\administrator**.
 - **InfrastructureAndExternal** Should be completed already.
 - **LocalAdmin** Select the SetupAdmin run as account.

- **Product Key** If you removed the data in Procedure 12, this line item will not appear. Also, toward the end of the VM deployment process, you will need to manually connect to the new VMs and select Skip to skip the product key entry during deployment. If you left the data present in Procedure 12, enter your Windows Server 2012 R2 product key in the format XXXXX-XXXXX-XXXXX-XXXXX-XXXXX.
- **VMClusterName** Enter **GWCLUS01**.
- **VMClusterStaticIPAddress** Enter **10.10.0.100**.

Setting	Value
DomainFQDN	contoso.com
DomainUserRAA	SetupAdmin
DomainUserRAAName	contoso\administrator
InfrastructureAndExternal	Management VM Network
LocalAdmin	SetupAdmin
VMClusterName	GWCLUS01
VMClusterStaticIPAddress	10.10.0.100

FIGURE 6-14 Configuring the final settings for the Windows Server Gateway service

14. With your settings completed, on the ribbon, click Refresh Preview to have the deployment service verify that your changes still result in appropriate host selection.
15. To view the settings that have been applied to your VMs, in the central canvas, click WINSERVERGW-VM1, and on the ribbon, click Properties.
 - Locations shows the path where the VM files will be stored.
 - Network Adapter 0 shows the IP information that will be applied to the VM at deployment time.
 - Machine Resources shows multiple virtual hard disks, all of which will be copied from the library to the target locations. These target locations have been determined by the use of the Infrastructure Storage classification earlier.

Click Cancel to close the WINSERVERGW-VM1 Properties window.
16. On the ribbon, click Deploy Service. In the Deploy Service dialog box, click Deploy.
17. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box. This may take several minutes to configure.

NOTE If you chose to use the evaluation mode instead of specifying a product key in Procedure 12, the deployment of the VM may seem to halt for a long time at 77 percent on the Customize Virtual Machine step. To overcome this, navigate to the VMs And Services workspace, expand All Hosts, and then GATEWAY. Expand GWCLUS, right-click GW01, and select Connect Via RDP. Enter your contoso\administrator credentials and log on to GW01. On GW01, open Hyper-V Manager, then double-click the WINSERVERGW-VM1 VM to open a console connection. When the console opens, you should see the VM is waiting for you to supply product key information. Click Skip, close the Virtual Machine Connection window, and return to VMM01. When you click Skip, the process should continue as expected.

18. While the first service is being deployed, minimize the Jobs window and navigate to the Library workspace. Expand Templates and click Service Templates. In the central pane, right-click the Windows Server 2012 R2 HA Gateway – 2NIC Service Template, and click Configure Deployment. The Select Name And Destination Window opens.
19. Repeat steps 6 through 17 of this procedure to create an additional Windows Server Gateway service with the following adjusted parameters. All other parameters should remain as documented in steps 3 through 16.
 - **Service Name** Windows Server Gateway 002
 - **VMClusterName** GWCLUS02
 - **VMClusterStaticIPAddress** 10.10.0.101
20. Upon successful completion of these steps, navigate to the VMs And Services workspace, expand the All Hosts Host Group, and select the GATEWAY Host Group.
21. On the ribbon, select Services. As shown in Figure 6-15, your new Windows Server Gateway 001 and Windows Server Gateway 002 services both indicate an OK status, and there is a single VM deployed successfully for each service and currently running.

Name	Status	All VMs Accessible	VM Status
Windows Server Gateway 001	OK	Yes	Running
Windows Server 2012 R2 Gateway	OK	Yes	Running
WINSERVERGW-VM1.contoso.com	Running		Running
Windows Server Gateway 002	OK	Yes	Running
Windows Server 2012 R2 Gateway	OK	Yes	Running
WINSERVERGW-VM2.contoso.com	Running		Running

FIGURE 6-15 The successful deployment of the two Windows Server Gateway services

With the services now deployed and operational, you can begin the integration of the Windows Server Gateway Hyper-V host, VMs, and System Center Virtual Machine Manager.

Procedure 13: Configure Windows Server Gateway 001 integration with System Center Virtual Machine Manager

In this procedure, you walk through integrating your newly deployed Windows Server Gateway service with System Center Virtual Machine Manager. By performing these actions, you bring the Windows Server Gateway under the management control of System Center Virtual Machine Manager. Subsequently, when new VM networks are deployed with network virtualization enabled, System Center Virtual Machine Manager will coordinate the management of these networks with the Windows Server Gateway.

It's important to note that the way you integrate the newly deployed service with System Center Virtual Machine Manager determines the scenario that the Windows Server Gateway enables. As discussed in Chapter 1, the Windows Server Gateway routes network traffic between the physical network and VM network resources, regardless of where the resources are located. You can use Windows Server Gateway to route network traffic between physical and virtual networks at the same physical location or at many different physical locations. For example, if you have both a physical network and a virtual network at the same physical location, as in this POC configuration, you can deploy a Hyper-V host that is configured with a Windows Server Gateway VM to act as a forwarding gateway and route traffic between the virtual and physical networks.

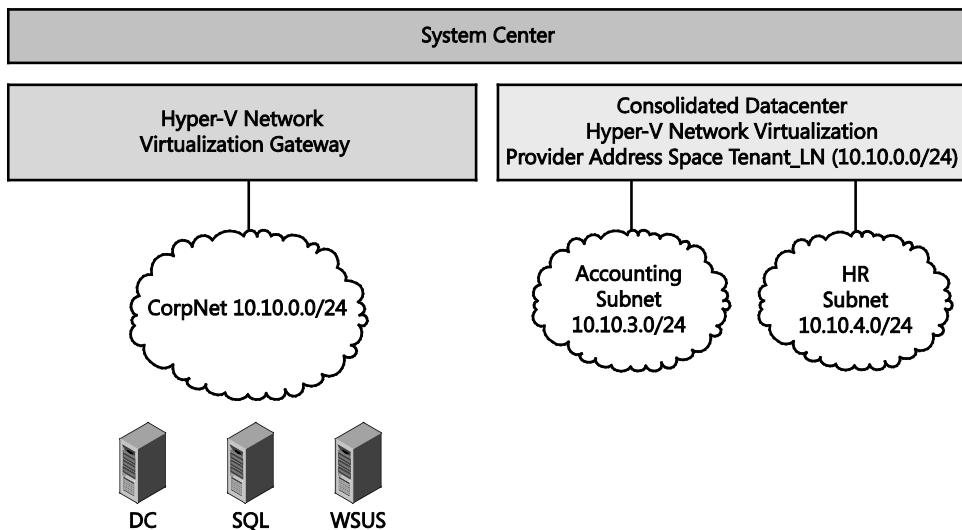


FIGURE 6-16 A representation of network virtualization configured with direct routing

Notice in Figure 6-16, on the left side, there are resources that are on a single network subnet called CorpNet. These workloads include your domain controllers, SQL Server instances, DNS services, and more. In your POC configuration, these are located on the 10.10.0.0/24 subnet. The Windows Server Gateway connects to this network and also communicates with System Center Virtual Machine Manager. On the right side of the graphic, you have several

virtual subnets. These virtual subnets are part of a single VM network. System Center Virtual Machine Manager assigns a Customer Address (CA) to each VM deployed into one of these virtual subnets. This is the IP address that is assigned by System Center Virtual Machine Manager into the VM. This address enables the VM to communicate with the rest of the network as if it had not been moved into a virtual network at all. The CA is visible to the VM and reachable from the non-virtual network. This assumes that you have the correct routing set up on your physical switches to route between the virtual subnets within the VM network.

You'll also notice that the consolidated datacenter box highlights the use of network virtualization but also references a Provider Address (PA) space. As discussed in Chapter 1, the PA address is the IP address that is assigned by System Center Virtual Machine Manager based on the underlying physical network infrastructure. In this POC configuration, the PA address space will come from the Tenant_LN logical network. The PA appears in the packets on the network that are exchanged with the server running Hyper-V that is hosting the VM. The PA is visible on the physical network but not to the VM itself.

This direct routing deployment scenario allows the enterprise to take advantage of Hyper-V network virtualization's ability to offer flexibility in both VM placement and cross-subnet live migration in the datacenter fabric. This increases datacenter efficiency, thereby reducing both operational expenditure (OPEX) and capital expenditure (CAPEX). Compare this with VLANs, where, for instance, allowing a VM to be migrated to a different server in a different rack may require additional administrative effort to add the specific VLANs to the top-of-rack and aggregation switches within those racks. All this effort is just to allow the VM to reside there. Network virtualization solves that problem.

In another example, if your virtual networks exist in the cloud, your cloud service provider (CSP) could deploy a Windows Server Gateway so that you can create a VPN site-to-site connection between your VPN server and the CSP's Windows Server Gateway. When this link is established, you can connect to your virtual resources in the cloud over the VPN connection.

In addition to the forwarding and site-to-site VPN already discussed, the Windows Server Gateway supports multi-tenant Network Address Translation (NAT) for VM Internet access. It also supports multi-tenant remote access VPN connections.

With this first deployment of the Windows Server Gateway, you will configure the integration between the Windows Server Gateway and System Center Virtual Machine Manager to specifically act as a forwarding gateway. For enterprises that deploy an on-premises private cloud, the Windows Server Gateway can act as a forwarding gateway and route traffic between virtual networks and the physical network. In this configuration, your Active Directory resources are currently located on a network that is not configured for network virtualization. This means that users within the environment, whether they are on laptops, desktops, or working with other physical servers, will be able to reach the domain controllers for authentication, name resolution, and other Active Directory-related services. Later, you will create some new VM networks that will have network virtualization enabled. You can create these VM networks for one or more of your departments, such as research and development or accounting. The Windows Server Gateway can route traffic between the virtual

network and the physical network so that those workloads can communicate back to workloads such as the domain controllers that are not using network virtualization.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the Fabric workspace.
2. Expand Networking, right-click Network Service and click Add Network Service. The Add Network Service Wizard appears.
3. On the Name page, in the Name text box, enter **Windows Server Gateway – Forwarder**, and then click Next.
4. On the Manufacturer And Model page, use the Model drop-down list to select Microsoft Windows Server Gateway, and click Next.
5. On the credentials page, click Browse, select SetupAdmin, click OK, and then click Next.
6. On the Connection String page, shown finished in Figure 6-17, enter the following information in the following format, and then click Next:

```
<Parameter1>=<Value1>;<Parameter2>=<Value2>;<ParameterN>=<ValueN>
```

- **VMHost=GWCLUS.contoso.com** This should be the fully qualified domain name of the Hyper-V host (or cluster, in a production environment). The hosts should be configured as dedicated Windows Server Gateway Hyper-V hosts.
- **GatewayVM=GWCLUS01.contoso.com** This should be the fully qualified domain name of the VM, or VM cluster that has been deployed in the Windows Server Gateway service.
- **BackendSwitch=Tenant_LN_LS** This should be the name of the virtual switch to automatically connect the back-end adapter for network virtualization. You only have a single logical switch deployed to your Windows Server Gateway Hyper-V host, so this should be used.
- **DirectRoutingMode=True** This confirms that this configuration will be for a forwarding gateway.
- **FrontEndServiceAddress=10.10.0.100** This is required if DirectRoutingMode is provided and set to True. Set this value as the IP address of this routing gateway. Network routing devices on the external network should point to this endpoint to get access to the VM network or networks behind the gateway.

```
Connection string:
VMHost=GWCLUS.contoso.com;GatewayVM=GWCLUS01.contoso.com;BackendSwitch=Tenant_LN_LS;DirectRoutingMode=True;FrontEndServiceAddress=10.10.0.100
```

FIGURE 6-17 The connection string to configure integration of the Windows Server Gateway

7. On the Certificates page, click Next.

8. On the Provider page, ensure that Microsoft Windows Server Gateway Provider is the selected configuration provider, and then click Test. A test results screen appears, as shown in Figure 6-18.

Test	Result
Connection API	Implemented
Test open connection	Passed
Capability discovery API	Implemented
Test capability discovery	Passed
Get certificate URL API	Implemented
Retrieve system info API	Implemented
Test system info	Passed
NAT management API	Implemented
Metering API	Implemented
Routing Domain Configuration API	Implemented
Customer Subnet Configuration API	Implemented


FIGURE 6-18 The successful testing of the Windows Server Gateway 001 service

9. If the test run is successful, click Next. On the Host Group page, select the All Hosts check box, and click Next.
10. On the Summary page, review the settings, and then click Finish. The Jobs window appears, showing the job status. Make sure that the job has a status of Completed, and then close the dialog box.
11. Return to the Fabric workspace, where a single network service is now listed. Right-click your Windows Server Gateway – Forwarder network service and select Properties. The Windows Server Gateway – Forwarder Properties window opens.
12. Configure System Center Virtual Machine Manager to recognize which network sites to use for the front-end and back-end adapters to ensure that it is able to allocate IP addresses from those sites and know which VM networks can use this Windows Server Gateway. Click the Connectivity tab.
13. Under Connectivity, notice that there are two check boxes. Select the Enable Front End Connection check box.
14. In this configuration, the front-end network is shared by management and also allows connectivity to a traditional physical network. In production environments, it is likely the front-end network will be separated. Use the FrontEnd Network Adapter drop-down list to select the adapter listed as contoso.com_10.10.0.0_24.
15. Use the Front End Network Site drop-down list to select Tenant_LN_0 (Tenant_LN).
16. Select the Enable Back End Connection check box, and then use the Back End Network Site drop-down list to select Tenant_LN_0 (Tenant_LN). This is the network that Hyper-V Network Virtualization uses to send encapsulated packets to and from the tenant VMs. To maintain isolation, in a production environment this network must be used only for Hyper-V Network Virtualization. In this POC configuration, the networks will be shared.
17. Click OK to close the Windows Server Gateway – Forwarder Properties window.

Procedure 14: Create a VM network and virtual subnets

With the Windows Server Gateway configured and integrated with System Center Virtual Machine Manager, you can create a VM network. In this procedure, you create a single VM network with multiple virtual subnets. The aim of this VM network is not to separate the contained workloads from services such as Active Directory, but to demonstrate how easy and flexible deployment of a new consolidated virtual network becomes when you use network virtualization. When you have completed the deployment of the new VM network, you will deploy several VMs into that VM network to check the connectivity.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace, and then click VM Networks.
2. Right-click VM Networks and select Create VM Network. The Create VM Network Wizard opens.
3. On the Name page, enter **CorpNet**. From the Logical Network drop-down list, select Tenant_LN since this is the logical network that has network virtualization enabled. Click Next.
4. On the Isolation page, ensure that Isolate Using Hyper-V Network Virtualization is selected. Note, that you can currently select either IPv4 or IPv6 for the protocols, not both. Click Next.
5. On the VM Subnets page, click Add. Under VM Subnet, enter **Accounting Subnet**, and for the Subnet, enter **10.10.3.0/24**.
6. Remaining on the VM Subnets page, click Add again, provide a name of **HR Subnet**, and for the Subnet, enter **10.10.4.0/24**. Click Next.
7. On the Connectivity page, integrate the new VM network with your recently deployed Windows Server Gateway. Under Connectivity, select the Connect Directly To An Additional Logical Network check box and ensure that the Direct Routing option is selected, as shown in Figure 6-19.



Connect directly to an additional logical network

- Direct routing
- Network address translation (NAT)

Select the gateway device to use for this connection.

Gateway device: Windows Server Gateway - Forwarder

Device capabilities: Direct routing

Remaining capacity: 1 VM networks
0 VPN tunnels

Front-end network: Front End (Tenant_LN_0)

FIGURE 6-19 Configuring the integration between a VM network and the Windows Server Gateway

8. On the Summary page, review your settings, and click Finish. The Jobs window appears to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.

9. Return to the VMs And Services workspace, ensure VM Networks is selected. Right-click the CorpNet VM network and select Create IP Pool. The Create Static IP Address Pool Wizard opens.
10. On the Name page, enter **Accounting IP Pool** for the name, and then click Next.
11. On the IP Address Range page, click Next.
12. On the Gateway page, for VMs within this VM network to be able to reach other workloads outside of the VM network, they will need a default gateway local to this subnet. Click Insert, and then click Enter Gateway Address. Enter **10.10.3.1**, and click Next.
13. On the DNS page, under Specify One Or More DNS Servers, click Insert, and then enter **10.10.0.12**. Click Insert again and enter **10.10.0.11**. Click Next.
14. On the WINS page, click Next.
15. On the Summary page, review your settings, and then click Finish. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
16. Repeat steps 9 through 15 of this procedure for an IP pool with the following details:
 - **Name** HR IP Pool
 - **VM Subnet** HR Subnet (10.10.4.0/24)
 - **Gateway** 10.10.4.1
 - **DNS** 10.10.0.12 and 10.10.0.11 (Ensure 10.10.0.11 is the top IP address is the list)

Procedure 15: Create accounting and HR VMs

With the VM network now deployed and the virtual subnets constructed, you can create and deploy a couple of VMs within this VM network to test the expected functionality and communication. To accelerate the deployment in this procedure, you use the System Center Virtual Machine Manager template capability. This is different from the more advanced service templates that you used to deploy the Windows Server Gateway services earlier.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the Library workspace, expand Templates, and click VM Templates.
2. On the ribbon, click Create VM Template. The Create VM Template Wizard opens.
3. On the Select Source page, ensure the Use An Existing VM Template Or A Virtual Hard Disk Stored In The Library option is selected, and click Browse. Select your WS2012R2 sysprepped virtual hard disk, click OK, and then click Next.
4. On the Identity page, for the VM Template name, enter **VM Template**, and click Next.
5. On the Configure Hardware page, under General, select Processor. Change the number of processors to 2.

6. Click Memory. Select the Dynamic option, and enter the following:
 - **Startup memory** 2048 MB
 - **Minimum memory** 1024 MB
 - **Maximum memory** 4096 MB
7. Under Bus Configuration, select the WS2012R2 virtual hard disk. Use the Classification drop-down list to select Tenant Storage.
8. Under Advanced, click Availability. Select the Make This Virtual Machine Highly Available check box, and then click Next.
9. On the Configure Operating System page, click Admin Password. Select the option to specify the password of the local administrator account, and enter your regular administrative password in both text boxes.
10. Select Product Key. If you are using a product key, enter it in the text box, and then click Next.
11. On the Application Configuration page, click Next.
12. On the SQL Server Configuration page, click Next.
13. On the Summary page, review your settings, and then click Create. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
14. Remaining in the VM Templates view, right-click your VM template and select Create Virtual Machine. The Create Virtual Machine Wizard opens.
15. On the Identity page, enter **ACT01** as the name, and then click Next.
16. On the Configure Hardware page, under Network Adapters, click Network Adapter 1.
17. Select the Connected To A VM Network option, and click Browse. Select CorpNet, and click OK.
18. Select Accounting Subnet from the VM Subnet drop-down list
19. Under IP Address, select the Static IP (From A Static IP Pool) option, and then click Next.
20. On the Configure Operating System page, click Next.
21. On the Select Destination page, use the Destination drop-down list to select COMPUTE, and click Next.
22. On the Select Host page, intelligent placement assesses the availability of each of your hosts and determines which is best suited for running this VM. Select the host with the highest ranking, and click Next.
23. On the Configure Settings page, review the settings. Click Network Adapter 0. Notice, under IPv4 Address From Logical Network, that System Center Virtual Machine Manager will use the Tenant_LN logical network to provide the Provider Address (PA) that was discussed earlier. Click Next.

24. On the Add Properties page, click Next.
25. On the Summary page, select the Start The VM After Deploying It check box, and then click Create. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
26. Repeat steps 14 through 25 of this procedure for a second VM, with the following changes from the previous steps:
 - **Virtual machine name** HR01
 - **Network Adapter 1** Connected to CorpNet, in the HR Subnet

Procedure 16: Test VM communication

With your VMs now deployed into the respective virtual subnets within the single CorpNet virtual network, it's important to ensure that the VMs can communicate with one another and to the rest of the network that is not using network virtualization. In this procedure, you review the communication between VMs and the rest of the network. In the next procedure, you'll review some of the inner workings of the network virtualization configuration to understand the changes that have taken place in the background.

In a physical network, a subnet is the Layer 2 (L2) domain where computers (virtual and physical) can directly communicate with each other without having to be routed. In Windows, if you statically configure a network adapter, you can set a default gateway that is the IP address to which you want to send all traffic that is going out of the particular subnet. This allows the traffic to be routed appropriately. This is typically the router for your physical network. To form a distributed router for a virtual network, Hyper-V network virtualization uses a built-in router that is part of every host. This means that every host, in particular the Hyper-V virtual switch, acts as the default gateway for all traffic that is going between virtual subnets (in this case, between the accounting subnet and HR subnet) that are part of the same VM network (in this case, CorpNet). In Windows Server 2012 R2, the address used as the default gateway is the lowest entry for the subnet (for example, it is the .1 address for a /24 subnet prefix). This address is reserved in each virtual subnet for the default gateway and cannot be used by VMs in the virtual subnet. You saw this when you defined the gateway for each of your IP pools for the virtual subnets you created earlier.

Acting as a distributed router, Hyper-V network virtualization provides an efficient way for all traffic inside a VM network to be routed appropriately because each host can directly route the traffic to the appropriate host without needing an intermediary. This is particularly true when two VMs in the same VM network but different virtual subnets are on the same physical host. The packet never has to leave the physical host.

It is important to note that for communication to work successfully between your virtual subnets (10.10.3.0/24 and 10.10.4.0/24) and any outside destination, you will have to configure your physical switch infrastructure to be able to route traffic into and out of this VM network and the respective virtual subnets. One example of an outside destination is your domain controller, which resides on a VM network that is not using network virtualization.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace.
2. Expand All Hosts, then expand COMPUTE and click on HVCLUSTER. There should be two VMs, ACT01 and HR01, in the central pane.
3. Right-click ACT01, click Connect Or View, and then click Connect Via Console to open a VM console session into ACT01.
4. Use the Ctrl-Alt-Del button at the top of the Virtual Machine Viewer window to enable the login. Enter your administrative password. At this point, your ACT01 VM is not joined to the domain.
5. Inside ACT01, right-click Start and select Command Prompt (Admin). The command prompt opens.
6. Enter **ipconfig**. As Figure 6-20 shows, for the Ethernet adapter, System Center Virtual Machine Manager has assigned the ACT01 VM an IP address on the 10.10.3.0/24 subnet. This is the CA that System Center Virtual Machine Manager has allocated from the accounting IP pool.

```

Connection-specific DNS Suffix . . . . . : 
Link-local IPv6 Address . . . . . : fe80::98c9:d23a:758b:9863%12
IPv4 Address. . . . . : 10.10.3.2
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 10.10.3.1

```

FIGURE 6-20 The result of an Ipconfig command on ACT01

7. Remaining within the command prompt, enter **firewall.cpl** to launch the firewall control panel. By default, the firewall blocks pings between operating systems. For test purposes and simplification, disable the firewall across all network types.
8. Minimize the Virtual Machine Viewer for ACT01. Repeat steps 3 through 7 of this procedure for HR01. You should find the HR01 VM has been assigned an IP address on the 10.10.4.0/24 subnet. This is the CA that System Center Virtual Machine Manager has allocated from the HR IP pool.
9. Remaining within the Virtual Machine Viewer for HR01, maximize your command prompt. Enter **ping 10.10.3.2**. You should receive a response. As mentioned earlier, Hyper-V is handling the routing between the virtual subnets within the CorpNet VM network.
10. With communication established between the HR and Accounting VMs, you can move on to test the communication back to the domain controllers and also to the Internet. Within HR01, click Start, and click on the Internet Explorer tile. Enter a website address of your choice. The website should resolve and display correctly. In this POC configuration, you are using the same edge default gateway as the rest of the organization is using, even though this VM resides in a virtualized network.
11. Return to the command prompt, ping 10.10.0.11 and then ping 10.10.0.12. Both of these domain controllers should respond. This demonstrates that you are able to communicate out of the virtualized network and onto the regular network.

12. Remaining within HR01, right-click the Start button, and click System.
13. In the System window, click Change settings.
14. In the System Properties window, click Change.
15. In the Computer Name/Domain Changes window, select the Domain option, enter **contoso.com**, and then click OK.
16. In the Windows Security window, enter **contoso\administrator** and your regular credentials, and then click OK.
17. When the Welcome To The Contoso.com Domain message appears, click OK. You have now successfully joined the domain from a VM that resides within a VM network, with traffic passing via the Windows Server Gateway to handle the encapsulation and de-encapsulation of the traffic. Click OK, click Close, and then click Restart Now to restart the machine.
18. When HR01 has completed rebooting, use the Ctrl-Alt-Del button at the top of the Virtual Machine Viewer window to enable the login. Click the back arrow, select Other User, and then enter your contoso\administrator credentials and password.
19. When you're logged in to HR01, right-click the Start button and select Command Prompt (Admin). The command prompt opens.
20. Ping DC01. You should now find that name resolution is working successfully and this machine, HR01, behaves just like any other on a non-virtualized network.

With HR01 configured, you can optionally repeat the steps in this procedure on ACT01. Your communication should now be working as expected.

Procedure 17: Review network virtualization configuration

With communication now operating as expected, it's important to take a minute to review what's happened in the background on both the GW01 dedicated Windows Server Gateway host and the WINSERVERGW-VM1 Windows Server Gateway VM. Both of these servers and System Center Virtual Machine Manager are playing a pivotal role in ensuring that communication among the different virtual subnets is working as expected—within and in and out of the CorpNet VM network.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace.
2. Expand All Hosts, then expand GATEWAY, and then expand GWCLUS. Right-click GW01 and select Connect Via RDP.
3. In the Windows Security window, enter your contoso\administrator credentials, and click OK to log in to GW01.
4. When you're logged in to GW01, on the taskbar, right-click the Windows PowerShell icon and select Run As Administrator.

5. In the Windows PowerShell window, enter **ipconfig**. There should be two IP addresses: one in the 10.10.1.0/24 subnet and one in the 10.10.2.0/24 subnet, both on the Datacenter_LN logical network.
6. Remaining in the Windows PowerShell window, enter **Get-NetVirtualizationProviderAddress**. This returns two provider addresses, as shown in Figure 6-21. (Your values may vary from the addresses displayed in Figure 6-21.) System Center Virtual Machine Manager allocated these PAs to GW01 from the Tenant_LN logical network. Although these have been allocated to the host, the PA addresses do not appear as regular IP addresses on the host; otherwise, they would have appeared when Ipconfig was running in the previous step. These IP addresses are used exclusively with network virtualization. System Center Virtual Machine Manager has allocated two PAs, one specifically for the host and another that is related to the cluster. In a production environment, where you have multiple Hyper-V hosts within the cluster, if a particular host were to fail, the PA that has been assigned as a clustered resource would fail over between hosts. This ensures that even with an outage of the underlying Hyper-V host, the Windows Server Gateway service is still operational and can still process traffic between the different networks.

```

ProviderAddress : 10.10.0.137
InterfaceIndex  : 23
PrefixLength    : 24
VlanID          : 0
AddressState    : Preferred
MACAddress      : d4ae528e8db8
ManagedByCluster : False

ProviderAddress : 10.10.0.136
InterfaceIndex  : 23
PrefixLength    : 24
VlanID          : 0
AddressState    : Preferred
MACAddress      : 001dd8b71c3c
ManagedByCluster : True

```

FIGURE 6-21 The result of running Get-NetVirtualizationProviderAddress on GW01

7. Remaining within the Windows PowerShell window, enter **Get-NetVirtualizationLookupRecord**. The Get-NetVirtualizationLookupRecord cmdlet gets lookup record policy entries for IP addresses that belong to a Hyper-V VM network. As mentioned earlier, network virtualization allows more than one virtual network to exist on the same physical network. Computers can use a CA within the virtual network to exchange network traffic with a VM. Network Virtualization manages the PAs that are the physical network addresses. This cmdlet returns records that map a CA to a PA. Within the results, scroll to locate the CA of 10.10.3.2. This is the IP address that was allocated inside the ACT01 VM. Notice that the record has a corresponding PA on the 10.10.0.0/24 subnet. This PA identifies the host that the VM ACT01 is currently running on. This PA was allocated to that particular Hyper-V host from the Tenant_LN logical network. Make a note of this PA. As shown in Figure 6-22, although you have the CA and PA, you also have more specific identifiers, such as the Virtual Subnet ID, which identifies which virtual subnet this VM is part of. Notice that

the 10.10.3.2 and 10.10.4.2 VMs have different Virtual Subnet IDs to represent the Accounting subnet and HR subnet, respectively. They do, however, have an identical CustomerID, which represents the VM network itself (CorpNet).

```
CustomerAddress : 10.10.3.2
VirtualSubnetID : 11449869
MACAddress      : 001dd8b71c34
ProviderAddress : 10.10.0.140
CustomerID     : {11D00839-A6AD-42DF-A1DA-379DEDDACD77}
Context        : SCVMM-MANAGED
Rule           : TranslationMethodEncap
VMName        :
UseVmMACAddress : False
Type          : Static

CustomerAddress : 10.10.4.2
VirtualSubnetID : 11218851
MACAddress      : 001dd8b71c3f
ProviderAddress : 10.10.0.139
CustomerID     : {11D00839-A6AD-42DF-A1DA-379DEDDACD77}
Context        : SCVMM-MANAGED
Rule           : TranslationMethodEncap
VMName        :
UseVmMACAddress : False
Type          : Static
```

FIGURE 6-22 The result of running `Get-NetVirtualizationLookupRecord` on GW01

8. To confirm that this information is correct, minimize your RDP session on GW01. Return to the System Center Virtual Machine Manager VMs And Services workspace, expand All Hosts, expand COMPUTE, and then expand HVCLUSTER.
9. In the central pane, confirm which host ACT01 is currently running on. Under HVCLUSTER, right-click the host that ACT01 is running on and select Connect Via RDP.
10. In the Windows Security window, enter your `contoso\administrator` credentials, and click OK to log in to the Hyper-V host.
11. When you're logged in to your Hyper-V host, on the taskbar, right-click the Windows PowerShell icon and select Run as Administrator.
12. In the Windows PowerShell window, enter **Install-WindowsFeature Hyper-V-PowerShell** to install the Hyper-V module for Windows PowerShell.
13. Type `Get-NetVirtualizationProviderAddress`, and press Enter. You should find that the PA listed here matches the result that was found on GW01. Minimize the RDP session to this Hyper-V host.
14. Return to the VMs And Services workspace, right-click the ACT01 VM and select Migrate Virtual Machine. The Migrate VM Wizard opens.
15. On the Select Host page, select a different Hyper-V host as the destination, ensure the Transfer Type is Live, and click Next.
16. On the Summary page, review the settings, and click Move.
17. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
18. Reopen your RDP session to GW01, and click on the Windows PowerShell window.

19. Rerun the `Get-NetVirtualizationLookupRecord` command. For the same CA (10.10.3.2) for ACT01, the corresponding PA has now changed because ACT01 now resides on a different Hyper-V host. If you were to migrate the ACT01 VM onto the same Hyper-V host as HR01, you would notice in the results that they would share the same PA, ensuring network virtualization can scale effectively.
20. Close the RDP session to GW01, and return to the System Center Virtual Machine Manager console. In the VMs And Services workspace, expand the GATEWAY Host Group, expand GWCLUS, and select GW01.
21. In the central pane, right-click `WINSERVERGW-VM1.contoso.com`, click `Connect Or View`, and then click `Connect Via Console`.
22. Use the `Ctrl-Alt-Del` button at the top of the Virtual Machine Viewer window to enable the login. Click the back arrow, select `Other User`, and then enter your `contoso\administrator` credentials and password.
23. When you're logged in to `WINSERVERGW-VM1`, on the taskbar, right-click the Windows PowerShell icon and select `Run As Administrator`.
24. In the Windows PowerShell window, enter `ipconfig`. Notice multiple IPv4 addresses are listed for your Ethernet adapter `contoso.com_10.10.0.0_24`. Every time you create a VM network and configure connectivity with the Gateway, an IP is assigned to that network and implemented on the Windows Server Gateway VM on the front-end network. Close the Virtual Machine Viewer window.
25. On `VMM01`, click `Start`, and type **Failover Cluster**. In the results on the right side, click `Failover Cluster Manager`.
26. When `Failover Cluster Manager` opens, in the top-left corner, right-click `Failover Cluster Manager` and select `Connect to Cluster`.
27. In the `Select Cluster` window, in the `Cluster Name` text box, enter **GWCLUS**, and click `OK`. This opens a connection to the underlying single-node Hyper-V Host cluster that the Windows Server Gateway VM is running on.
28. Expand `GWCLUS.contoso.com`, and click `Roles`. As Figure 6-23 shows, in the central pane, a role has been configured and is running. When you integrated the Windows Server Gateway with System Center Virtual Machine Manager, this floating PA was configured. You'll notice that the `10.10.0.136` figure matches one of the PAs you saw earlier when you ran the command `Get-NetVirtualizationProviderAddress` on `GW01`. This is the PA that would float between physical cluster nodes and fail over in the event of an outage.

Roles (1)					
Search					
Name	Status	Type	Owner Node	Priority	Information
HNV PA 10.10.0.136	Running	Other	GW01	Medium	

FIGURE 6-23 The network virtualization role service that is running on the Hyper-V cluster

29. Repeat steps 26 through 28 of this procedure for GWCLUS01. This is the guest cluster, currently with a single node, that is running within WINSERVERGW-VM1. In the Roles view, a role called Hyper-V network virtualization Gateway is currently running. In a production environment with a two-node guest cluster, the role would fail over to the other corresponding virtual cluster node should one of the virtual cluster nodes fail.
30. In the bottom pane, click the Resources tab. A number of resources are associated with this running role. These role services are the key IP addresses associated with the role discussed above that ensure communication through the Windows Server Gateway. Close Failover Cluster Manager, any Virtual Machine Viewer window, and all RDP sessions that are currently established from VMM01.

Procedure 18: Configure Windows Server Gateway 002 integration with System Center Virtual Machine Manager

With the CorpNet network up and running, and communication between the respective virtual subnets functioning as expected, you can implement an alternative network virtualization scenario deployment. The CorpNet network virtualization scenario was focused on direct routing, in which the Windows Server Gateway that you configured was used exclusively by the CorpNet VM network but no others.

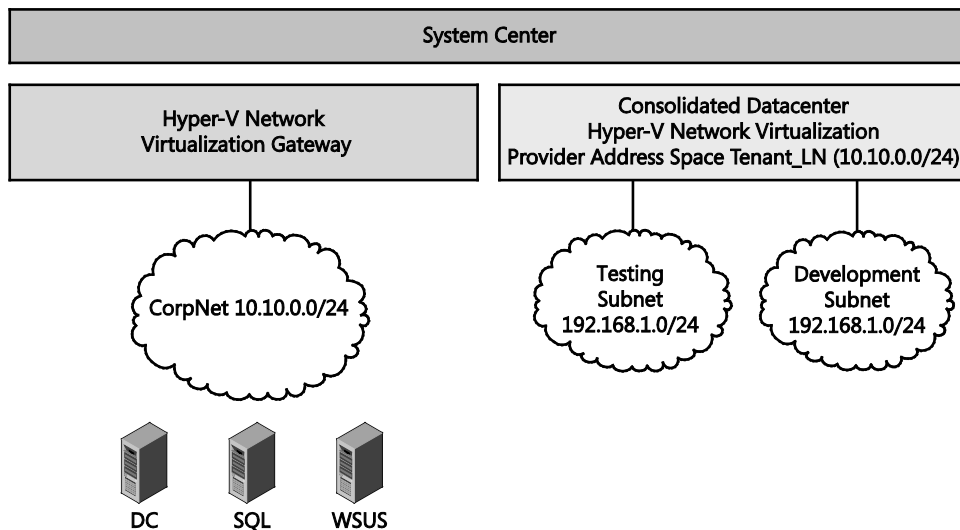


FIGURE 6-24 A representation of network virtualization configured with NAT and multiple, isolated VM networks

In this next scenario, illustrated in Figure 6-24, you walk through the configuration of the other Windows Server Gateway service that was successfully deployed earlier. However, this time, you will configure the integration between Windows Server Gateway 002 and System Center Virtual Machine Manager to support a multitenant NAT scenario. This scenario applies

to enterprises that want to segregate departmental workloads (perhaps for the test and development teams) from the main CorpNet network. Alternatively, this could be a service provider that wants to host VMs from many different customers that obviously require stringent isolation but also require communication out to the Internet.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the Fabric workspace.
2. Expand Networking, right-click Network Service and select Add Network Service. The Add Network Service Wizard opens.
3. On the Name page, in the Name text box, enter **Windows Server Gateway – NAT**, and click Next.
4. On the Manufacturer And Model page, use the Model drop-down list to select Microsoft Windows Server Gateway, and click Next.
5. On the Credentials page, click Browse, select SetupAdmin, click OK, and then click Next.
6. On the Connection String page, enter the following information in the format
`<Parameter1>=<Value1>;<Parameter2>=<Value2>;<ParameterN>=<ValueN>`
 - **VMHost=GWCLUS.contoso.com** This is the fully qualified domain name of the Hyper-V host (or cluster, in a production environment). The hosts should be configured as dedicated Windows Server Gateway Hyper-V hosts.
 - **GatewayVM=GWCLUS02.contoso.com** This is the fully qualified domain name of the VM or VM cluster that has been deployed in the Windows Server Gateway service.
 - **BackendSwitch=Tenant_LN_LS** This is the name of the virtual switch to automatically connect the back-end adapter for network virtualization. You only have a single logical switch deployed to your Windows Server Gateway Hyper-V host, so this should be used.

Click Next. Figure 6-25 shows the completed connection string.




FIGURE 6-25 The connection string to configure integration of the Windows Server Gateway

7. On the Certificates page, click Next.
8. On the Provider page, ensure that Microsoft Windows Server Gateway Provider is the selected configuration provider, and then click Test.
9. After a successful test run, click Next. On the Host Group page, select the All Hosts check box, and click Next.

10. On the Summary page, review the settings, and then click Finish. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
11. Return to the Fabric workspace where two network services should now be listed. Right-click your Windows Server Gateway – NAT network service and select Properties. The Windows Server Gateway – NAT Properties window opens.
12. Configure System Center Virtual Machine Manager to recognize which network sites to use for the front-end and back-end adapters to ensure that it is able to allocate IP addresses from those sites and know which VM networks can use this Windows Server Gateway. Click the Connectivity tab.
13. Under Connectivity, notice two check boxes. Select the Enable Front End Connection check box.
14. In this configuration, the front-end network is shared by management and also allows connectivity to a traditional physical network. In production environments, the front-end network will likely be separated. Use the Front End Network Adapter drop-down list to select the adapter listed as contoso.com_10.10.0.0_24.
15. Use the Front End Network Site drop-down list to select Tenant_LN_0 (Tenant_LN).
16. Select the Enable Back End Connection check box, and then use the Back End Network Site drop-down list to select Tenant_LN_0 (Tenant_LN). This is the network that Hyper-V network virtualization uses to send encapsulated packets to and from the tenant VMs. To maintain isolation in a production environment, this network must only be used for Hyper-V network virtualization. In this POC configuration, the networks will be shared.
17. Click OK to close the Windows Server Gateway – NAT Properties window.

Procedure 19: Create VM networks and virtual subnets

With the Windows Server Gateway 002 configured and integrated with System Center Virtual Machine Manager, you can create two VM networks. In this procedure, you create a VM network for the Testing team, with a single virtual subnet, and another VM network for the Development team, again with a single virtual subnet. The aim of these VM networks is to be isolated from the main CorpNet virtual network that was deployed earlier and also to be isolated from the other networks in the environment. These include the 10.10.0.0/24 network that your domain controllers, WDS, WSUS, and SQL Server VMs reside on, and also the 10.10.1.0/24 and 10.10.2.0/24 networks that the various hosts reside on. When you have completed the deployment of the new VM network, you will deploy several VMs into that VM network to check the connectivity.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace, and then click VM Networks.
2. Right-click VM Networks and select Create VM Network. The Create VM Network Wizard opens.

3. On the Name page, enter **Testing**. From the Logical Network drop-down list, select Tenant_LN. This is the logical network that has network virtualization enabled. Click Next.
4. On the Isolation page, ensure that Isolate Using Hyper-V Network Virtualization is selected. Note that you can currently select either IPv4 or IPv6 for the protocols, not both. Click Next.
5. On the VM Subnets page, click Add. Under VM Subnet, enter a name of **Testing Subnet**, and for Subnet, enter **192.168.1.0/24**.
6. On the Connectivity page, you will integrate the new VM network with your recently deployed Windows Server Gateway. Under Connectivity, select the Connect Directly To An Additional Logical Network check box, and ensure the Network Address Translation (NAT) option is selected. Notice that each Windows Server Gateway service that is configured for NAT supports up to 50 different VM networks. Click Next.
7. On the Network Address Translation (NAT) page, you can choose to configure a specific IP that will be used as the single NAT IP address, or you can let System Center Virtual Machine Manager deploy one from the Tenant_LN logical network IP pool. For now, leave the NAT Rules box empty, and click Next.
8. On the Summary page, review your settings, and click Finish. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
9. Return to the VMs And Services workspace, and ensure VM Networks is selected. Right-click the Testing VM network and select Create IP Pool. The Create Static IP Address Pool Wizard opens.
10. On the Name page, enter **Testing IP Pool** for the name, and then click Next.
11. On the IP Address Range page, click Next.
12. On the Gateway page, you can leave the gateway IP option empty, or you can specify 192.168.1.1 since this will be the default by design. System Center Virtual Machine Manager automatically uses the x.x.x.1 address for its default gateway within a VM network, and this cannot be modified. Click Next.
13. On the DNS page, since this VM network will be isolated from your Active Directory infrastructure, select DNS servers that can resolve public website addresses. In this example, you could use 208.67.222.222 and 208.67.220.220 or specify your own public DNS servers of choice. Click Next.
14. On the WINS page, click Next.
15. On the Summary page, review your settings, and then click Finish. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box

16. Repeat steps 2 through 15 of this procedure for a VM Network and IP pool with the following details:
- **VM Network Name** Development
 - **VM Subnet** Development Subnet and 192.168.1.0/24
 - **Connectivity** NAT
 - **IP Pool Name** Development IP Pool
 - **VM Subnet** Development Subnet (192.168.1.0/24)
 - **DNS** 208.67.222.222 and 208.67.220.220

Procedure 20: Create testing and development VMs

The VM networks are now deployed, and the virtual subnets and IP pools are constructed. You can now create and deploy a couple of VMs across the two VM networks to test the expected functionality and communication. To accelerate the deployment in this procedure, you use the template that was created earlier.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the Library workspace, expand Templates, and click VM Templates.
2. Right-click your VM template and select Create Virtual Machine. The Create Virtual Machine Wizard opens.
3. On the Identity page, enter **TEST01** as the name, and then click Next.
4. On the Configure Hardware page, under Network Adapters, click Network Adapter 1.
5. Select the Connected To A VM Network option, and click Browse. Select Testing, and click OK.
6. Select Testing Subnet from the VM Subnet drop-down list.
7. Under IP address, select the Static IP (from a static IP pool) option, and then click Next.
8. On the Configure Operating System page, click Next.
9. On the Select Destination page, use the Destination drop-down list to select COMPUTE, and click Next.
10. On the Select Host page, intelligent placement will assess the availability of each of your hosts and determine which is best suited for running this VM. Select the host with the highest ranking, and click Next.
11. On the Configure Settings page, review the settings. Click Network Adapter 0. Notice under IPv4 Address From Logical Network that System Center Virtual Machine Manager will use the Tenant_LN logical network to provide the PA that was discussed earlier. Click Next.
12. On the Add Properties page, click Next.

13. On the Summary page, select the Start The VM After Deploying It check box, and then click Create. The Jobs window opens to show the job status. Make sure that the job has a status of Completed, and then close the dialog box.
14. Repeat steps 2 through 13 of this procedure for a second VM, making the following changes:
 - **Virtual machine name** DEV01
 - **Network Adapter 1** Connected to Development, in the Development Subnet

Procedure 21: Test VM communication

With your VMs now deployed into the respective virtual networks, it's important to ensure that the VMs can communicate to the outside world. Because your configuration uses NAT, the VMs will automatically be granted Internet connectivity using the Windows Server Gateway. The VMs will use the public DNS servers that you specified during the creation of the VM network and IP pool.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace.
2. Expand All Hosts, expand COMPUTE, and click HVCLUSTER. Your new VMs, DEV01 and TEST01, should be in the central pane.
3. Right-click DEV01, select Connect Or View, and then click Connect Via Console to open a VM console session into DEV01.
4. Use the Ctrl-Alt-Del button at the top of the Virtual Machine Viewer window to enable the login. Enter your administrative password. Your DEV01 VM is now joined to the domain.
5. Inside DEV01, right-click the Start button and select Command Prompt (Admin). The command prompt opens.
6. Enter **ipconfig**. For the Ethernet adapter, System Center Virtual Machine Manager has assigned the DEV01 VM an IP address on the 192.168.1.0/24 subnet. This is the CA that System Center Virtual Machine Manager has allocated from the Development IP pool. Because this is the first VM deployed in the pool, it most likely has an IP address of 192.168.1.2 and a default gateway of 192.168.1.1, even though you did not specify a default gateway when creating the IP pool for this VM network.
7. Minimize the Virtual Machine Viewer for DEV01. Repeat steps 3 through 7 of this procedure for TEST01. You should find the TEST01 VM has been assigned an IP address on the 192.168.1.0/24 subnet. Even though it will have been assigned the 192.168.1.2 IP address, which is the same as DEV01, the two VMs will not conflict because they are isolated within different VM networks. This allows IT to create many VM networks and use many identical subnets. From a service provider perspective, it also allows the service provider's customers to bring their own IP address into the service provider hosting environment, safe in the knowledge that their IP scheme will not overlap or conflict with that of another customer.

8. Remaining within the Virtual Machine Viewer for TEST01, click Start, and click the Internet Explorer tile. Enter a website address of your choice. The website should resolve correctly and display. You are using the public DNS servers specified in the creation of the Testing IP pool to resolve the addresses.
9. To confirm the use of the public DNS servers, within your Windows PowerShell window, launch nslookup. Your default server and the IP address should be listed. Enter **microsoft.com**. The lookup should be successful. Minimize all Virtual Machine Viewer windows.

Procedure 22: Configure NAT rules for VM networks

With your VMs now deployed into the respective virtual networks and connectivity to the outside world confirmed, it's important to ensure that the testing and development teams can communicate into their respective VM. With the current configuration, however, the two VMs that have just been deployed are completely isolated from the other networks within the environment. This ensures that the VMs themselves can't reach the sensitive management and authentication services on the management network. The challenge is therefore allowing the development and testing teams to reach their respective VMs to perform their jobs. Their local workstations won't be able to reach inside the VM networks like they could in the previous direct routing configuration.

This is where NAT rules come in. A NAT rule defines an entry point into the VM network to a specific VM, over a specific port, from outside of that virtual network. From there, the development and testing teams can hop to their other respective VMs within their VM networks. Alternatively, imagine a scenario where a service provider was hosting VMs that were serving public-facing websites. A NAT rule can enable people on the Internet to reach that website that is running on a specific VM in an isolated VM network for a particular protocol and port.

In this procedure, you create two NAT rules to ensure that your respective development and testing teams can use Remote Desktop to reach their VMs from outside the respective VM networks.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace, and then click VM Networks.
2. Right-click the Development VM network and select Properties.
3. Click the Network Address Translation (NAT) tab, and then click Add. Figure 6-26 shows the completed NAT rule. Enter the following information for the NAT rule. Take note of the IP address listed under the Tenant_LN_Pool (10.10.0.100 – 10.10.0.250). This is the IP address that System Center Virtual Machine Manager has allocated as the NAT IP address for this VM network.
 - **Name** DEV01 RDP
 - **Protocol** TCP
 - **Incoming Port** 3389

- **Destination IP** 192.168.1.2
- **Destination Port** 3389

IP address pool: Tenant_LN_Pool (10.10.0.100 - 10.10.0.250)

IP address: 10.10.0.145

Specify network address translation (NAT) rules:

Name	Protocol	Incoming Port	Destination IP	Destination P...	Add
DEV01 RDP	TCP	3389	192.168.1.2	3389	Remove

FIGURE 6-26 The NAT rule for the development VM network

4. Repeat steps 2 and 3 of this procedure for TEST01. This adds RDP as a rule in the Testing VM network. All other settings should be the same.
5. Open your Virtual Machine Viewer connection to DEV01. Inside the VM, right-click Start and select System.
6. Within the System window, click Remote Settings. In the System Properties window, select the Allow Remote Connections To This Computer option, and click OK.
7. Repeat steps 5 and 6 of this procedure on TEST01.
8. On VMM01, click Start, and enter **mstsc**. Click Remote Desktop Connection in the results.
9. In the Remote Desktop Connection window, in the Computer text box, enter the NAT IP address of the development VM network. It is important you enter this IP address instead of the VM IP address. VMM01 cannot route to the 192.168.1.0/24 network because it is isolated through network virtualization. Because you're using the NAT IP address and the NAT rule, RDP traffic on port 3389 will be passed through to the 192.168.1.0/24 network, specifically to the DEV01 VM, with the 192.168.1.2 IP address. Click Connect.
10. In the Windows Security window, enter your administrative credentials, and click OK. You should successfully log in to the VM, demonstrating a correctly configured and working network virtualization with NAT configuration. Close the RDP session.
11. Repeat steps 9 and 10 of this procedure with the Testing VM network and TEST01 VM.

It's important to note that this NAT rule configuration would only allow the development and testing teams to use Remote Desktop to connect directly to the VM that was specified in the NAT rule. You could add additional rules with different ports that pointed to a different VM within the VM network. The development or testing users could then use an IP address of 10.10.0.145:3390, and the rule would forward them to an alternative VM, perhaps one with an IP address of 192.168.1.3, inside the VM network.

Alternatively, the development and testing teams could use the single VM as a jump point to reach the other VMs within their virtual network. You could also establish a Windows Server Remote Desktop Gateway within the VM network to direct users to different VMs after they have entered the VM network.

A more production-ready option would be to use a new capability known as the Remote Console in System Center 2012 R2. You can find more information on this can on TechNet at <http://technet.microsoft.com/en-us/library/dn469415.aspx>.

Procedure 23: Review network virtualization configuration

With NAT communication now operating as expected, it's important to take a minute to review what's happened in the background on both the GW01 dedicated Windows Server Gateway host and the WINSERVERGW-VM2 Windows Server Gateway VM. Both of these servers and System Center Virtual Machine Manager are playing a pivotal role in ensuring that communication in and out of the VM networks is working as expected.

1. On VMM01, in the System Center Virtual Machine Manager console, navigate to the VMs And Services workspace.
2. Expand All Hosts, then expand GATEWAY, and then expand GWCLUS. Right-click GW01 and select Connect Via RDP.
3. In the Windows Security window, enter your contoso\administrator credentials, and click OK to log in to GW01.
4. When you're logged in to GW01, on the taskbar, right-click the Windows PowerShell icon and select Run As Administrator.
5. In the Windows PowerShell window, enter **Get-NetVirtualizationLookupRecord -CustomerAddress 192.168.1.2**. The Get-NetVirtualizationLookupRecord cmdlet gets lookup record policy entries for IP addresses that belong to a Hyper-V VM network. By explicitly including the -CustomerAddress parameter of 192.168.1.2, you can scope the results down to those related to the new VM networks for testing and development. As shown in Figure 6-27, the VMs have identical CAs yet can be distinguished by different VirtualSubnetIDs, CustomerIDs, and also by the PA, which indicates these virtual machines are currently residing on different hosts.

```
CustomerAddress : 192.168.1.2
VirtualSubnetID : 16344052
MACAddress      : 001dd8b71c43
ProviderAddress : 10.10.0.146
CustomerID      : {3865E9CC-AF1F-42BE-B2A4-6ED0BC36F703}
Context         : SCVMM-MANAGED
Rule            : TranslationMethodEncap
VMName          :
UseVmMACAddress : False
Type            : Static

CustomerAddress : 192.168.1.2
VirtualSubnetID : 4233631
MACAddress      : 001dd8b71c45
ProviderAddress : 10.10.0.147
CustomerID      : {D8026EF8-AA50-495F-8973-1523CA82DA2D}
Context         : SCVMM-MANAGED
Rule            : TranslationMethodEncap
VMName          :
UseVmMACAddress : False
Type            : Static
```

FIGURE 6-27 The result of running Get-NetVirtualizationLookupRecord -CustomerAddress 192.168.1.2 on GW01

6. To confirm that this information is correct, minimize your RDP session on GW01. Return to the System Center Virtual Machine Manager VMs And Services workspace, expand All Hosts, expand COMPUTE, and then expand HVCLUSTER.
7. In the central pane, confirm which host DEV01 is currently running on. Under HVCLUSTER, right-click the host that DEV01 is running on and select Connect Via RDP.
8. In the Windows Security window, enter your contoso\administrator credentials, and click OK to log in to the Hyper-V host.
9. When you're logged in to your Hyper-V host, on the taskbar, right-click the Windows PowerShell icon and select Run As Administrator.
10. In the Windows PowerShell window, enter **Install-WindowsFeature Hyper-V-PowerShell** to install the Hyper-V module for Windows PowerShell.
11. Remaining in the Windows PowerShell window, enter **Get-NetVirtualizationProviderAddress**. The PA listed here should match the result found on GW01.
12. Close the RDP session to GW01 and return to the System Center Virtual Machine Manager console. In the VMs And Services workspace, expand the GATEWAY Host Group, expand GWCLUS, and select GW01.
13. In the central pane, right-click WINSERVERGW-VM2.contoso.com, click Connect Or View, and then click Connect Via Console.
14. Use the Ctrl-Alt-Del button at the top of the Virtual Machine Viewer window to enable the login. Click the back arrow, select Other User, and then enter your contoso\administrator credentials and password.
15. When you're logged in to WINSERVERGW-VM2, on the taskbar, right-click the Windows PowerShell icon and select Run As Administrator.
16. In the Windows PowerShell window, enter **ipconfig**. Notice that for your Ethernet adapter, contoso.com_10.10.0.0_24 lists multiple IPv4 addresses. Every time you create a VM network and configure connectivity with the Gateway, an IP is assigned to that network and implemented on the Windows Server Gateway VM on the front-end network.
17. Remaining in the Windows PowerShell window, enter **Get-NetNatExternalAddress**. This displays the current IP addresses used for NAT for the development and testing VM networks.
18. In the Windows PowerShell window, enter **Get-NetNatSession**. Review the information presented on currently in-use NAT sessions, such as those using the RDP rule that you defined earlier. Close the Virtual Machine Viewer window.
19. On VMM01, click Start, and type **Failover Cluster**. In the results on the right side, click Failover Cluster Manager.
20. When Failover Cluster Manager opens, in the top-left corner, right-click Failover Cluster Manager and select Connect to Cluster.

21. In the Select Cluster window, in the Cluster Name text box, enter **GWCLUS02**, and click OK. This is the guest cluster, currently with a single node, that is running within WINSERVERGW-VM2. The Roles view shows a role called Hyper-V Network Virtualization Gateway is currently running. In a production environment with a two-node guest cluster, if one of the virtual cluster nodes should fail, the role would fail over to the other corresponding virtual cluster node.
22. In the bottom pane, click the Resources tab. Several resources are associated with this running role. More are associated with this NAT configuration than you previously saw with the direct routing example. Close Failover Cluster Manager, any Virtual Machine Viewer windows, and all RDP sessions that are currently established from VMM01.

That concludes this chapter and the book. Throughout this book, you've learned about how to deploy key management, compute, storage, and networking capabilities. You started from bare-metal and built up to a fully virtualized, software-defined infrastructure. Now you can continue to deploy and test your own key workloads, evaluate other key Hyper-V and System Center capabilities, and continue learning about the Microsoft software-defined datacenter technologies.

About the authors



MATT MCSPIRIT is a Senior Technical Product Marketing Manager in the Cloud + Enterprise business group, with a focus on Virtualization and Business Continuity. A Microsoft employee since January 2006, Matt's also an MCSE, MCITP: Virtualization Administrator, and a VMware VCP5-DCV with extensive experience across a broad portfolio of both Microsoft and non-Microsoft technologies. Matt has presented at a number of high-profile events, including TechEd, MMS, the launches of Windows Server 2008 through to Windows Server 2012 R2, Private Cloud Roadshows, and many more.



JIM KELLEY has been a systems engineer for 20 years, specializing in enterprise platforms, application services, cloud infrastructures, System Center, SQL Server, and DW/BI. Over the years, Jim has maintained several certifications including MSCE, MCITP, MSDBA, MCT, and CCNA. He has been a Microsoft contractor for 15 years, working with several groups including Premiere Support, Windows Server Clustering, Global Technical Readiness, Active Directory, Platforms, SMB partner readiness, and SQL Server teams. His specialties include Windows Server architecture and deployment, management, high availability, networking, security governance, and various application platforms like SharePoint, IIS, and SQL Server. Jim also works with veterans groups to help retrain returning veterans to prepare them to re-enter the workforce.

This page intentionally left blank

About the series editor



MITCH TULLOCH is a well-known expert on Windows Server administration and cloud computing technologies. He has published hundreds of articles on a wide variety of technology sites and has written, contributed to or been series editor for over 50 books. Mitch is one of the most popular authors at Microsoft Press—the almost two dozen ebooks on Windows Server and System Center he either wrote or was Series Editor on have been downloaded more than 2.5 million times! For a complete list of

free ebooks from Microsoft Press, visit the Microsoft Virtual Academy at <http://www.microsoftvirtualacademy.com/ebooks>.

Mitch has repeatedly received Microsoft's Most Valuable Professional (MVP) award for his outstanding contributions to supporting the global IT community. He is a ten-time MVP in the technology area of Windows Server Software Packaging, Deployment & Servicing. You can find his MVP Profile page at <http://mvp.microsoft.com/en-us/mvp/Mitch%20Tulloch-21182>.

Mitch is also Senior Editor of WServerNews, a weekly newsletter focused on system admin and security issues for the Windows Server platform. With almost 100,000 IT pro subscribers worldwide, WServerNews is the most popular Windows Server–focused newsletter in the world. Visit <http://www.wservernews.com> and subscribe to WServerNews today!

Mitch also runs an IT content development business based in Winnipeg, Canada, that produces white papers and other collateral for the business decision maker (BDM) and technical decision maker (TDM) audiences. His published content ranges from white papers about Microsoft cloud technologies to reviews of third-party products designed for the Windows Server platform. Before starting his own business in 1998, Mitch worked as a Microsoft Certified Trainer (MCT) for Productivity Point.

For more information about Mitch, visit his website at <http://www.mtit.com>. You can also follow Mitch on Twitter @mitchtulloch.

This page intentionally left blank



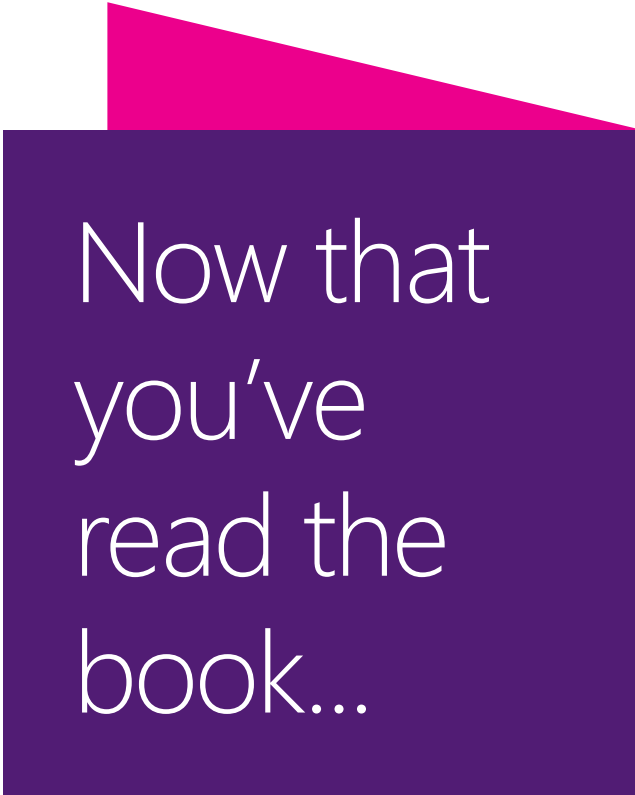
From technical overviews to drilldowns on special topics, get *free* ebooks from Microsoft Press at:

www.microsoftvirtualacademy.com/ebooks

Download your free ebooks in PDF, EPUB, and/or Mobi for Kindle formats.

Look for other great resources at Microsoft Virtual Academy, where you can learn new skills and help advance your career with free Microsoft training delivered by experts.

Microsoft Press



Now that
you've
read the
book...

Tell us what you think!

Was it useful?

Did it teach you what you wanted to learn?

Was there room for improvement?

Let us know at <http://aka.ms/tellpress>

Your feedback goes directly to the staff at Microsoft Press,
and we read every one of your responses. Thanks in advance!

