



Introducing Windows Server 2016 Technical Preview

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

Copyright © 2016 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: 2014958507

ISBN: 978-0-7356-9773-7

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/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> 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: Dianne Russell, Octal Publishing, Inc.

Copyeditor: Bob Russell, Octal Publishing, Inc.

Cover: Twist Creative • Seattle

Visit us today at

microsoftpressstore.com

- **Hundreds of titles available** – Books, eBooks, and online resources from industry experts
- **Free U.S. shipping**
- **eBooks in multiple formats** – Read on your computer, tablet, mobile device, or e-reader
- **Print & eBook Best Value Packs**
- **eBook Deal of the Week** – Save up to 60% on featured titles
- **Newsletter and special offers** – Be the first to hear about new releases, specials, and more
- **Register your book** – Get additional benefits



Contents

Introduction.....	vi
Acknowledgments.....	vi
Free ebooks from Microsoft Press	vii
Errata, updates, & book support	vii
We want to hear from you.....	viii
Stay in touch.....	viii
Chapter 1: The software-defined datacenter.....	1
Delivering agile IT	2
Understanding the Microsoft Cloud Platform.....	3
Facilitating enterprise mobility	3
Creating the Internet of Your Things	4
Providing the right environment for application innovation.....	4
Unlocking insights on any data.....	4
Transforming the datacenter.....	4
How to transform your datacenter	5
Building a software-defined foundation	6
What makes it possible.....	7
Automating and securing your infrastructure.....	10
Extending to the cloud on demand.....	11
Chapter 2: Compute.....	15
Failover Clustering improvements	15
Creating a cloud witness by using Azure.....	15
Shared VHDX improvements.....	18
Improved cluster logs.....	20
Active memory dump.....	22
Network name diagnostics.....	23
Cluster Operating System Rolling Upgrade	24

Hyper-V improvements.....	30
VM groups	30
True VM mobility.....	35
VM configuration version	40
New configuration file format.....	41
Production checkpoints.....	43
Hot add and hot remove for network adapters and memory	44
MultiPoint Services.....	48
Windows MultiPoint Server.....	48
Scenario for better understanding.....	50
MultiPoint Services role in Windows Server 2016 Technical Preview	50
MultiPoint Manager and MultiPoint Dashboard	57
Chapter 3: Storage	63
Storage Replica.....	63
Synchronous replication.....	64
Asynchronous replication	64
Implementation-specific details.....	65
Requirements.....	66
Recommendations.....	66
Scenarios	67
Storage Replica in Windows Server 2016 Technical Preview.....	70
Storage Spaces Direct.....	72
Implementation details.....	73
Improved scalability	74
Storage Spaces optimized pool	74
Failure scenarios	74
Learn more.....	76
Deduplication	76
Storage Quality of Service.....	78
Chapter 4: Networking	81
Software-defined network	81
Network virtualization	81
Network Controller.....	84
RAS Gateway Multitenant BGP router	86
Software Load Balancing.....	87
Datacenter firewall.....	88

Web Application Proxy	89
Publishing capability enhancements	89
Publishing Exchange Server 2013	93
Web Application Proxy troubleshooting	100
Collecting information about your environment	100
Using the Microsoft Exchange Best Practices Analyzer	100
Certificate issues	102
Configuration data in AD FS is inconsistent or corrupt	102
Supporting non-SPI-capable clients	103
Chapter 5: Security	104
Shielded VMs	104
Threat-resistant technologies	106
Control Flow Guard	106
Device Guard (Code Integrity)	106
Credential Guard	109
Windows Defender	112
Threat detection technologies	112
Securing privileged access	114
Just In Time and Just Enough Administration	114
A strategy for securing privileged access	116
Short-term plan	117
Medium-term plan	118
Long-term plan	119
Chapter 6: App Plat	121
Nano Server	121
Understanding Nano Server	121
Deploying Nano Server	124
Specializing Nano Server	125
Remotely managing Nano Server	125
Containers	127
What is a container?	127
Why use containers?	129
Windows Server containers versus Hyper-V containers	129
What about Docker?	133

Chapter 7: Systems management.....	136
Windows PowerShell improvements	136
DSC Local Configuration Manager.....	136
New methods in LCM.....	142
DSC partial configurations.....	143
Setting up the LCM Meta Configuration	144
Authoring the configurations.....	145
Deploying the configurations	147
PowershellGet and NuGet	148
System Center 2016.....	151
Operations Management Suite.....	154
About the author	161
About the contributors	162

Introduction

Windows Server has powered a generation of organizations, from small businesses to large global enterprises. No matter what your role in IT, you can be guaranteed that you have touched Windows Server at some point in your career or, at the very least, you have seen it from afar! No matter what your area of expertise, this ebook introduces you to Windows Server 2016 Technical Preview and its latest developments, which is the next version of Windows Server.

Each chapter has been written by either field experts or members of the product group, who provide you with the latest information on every improvement or new feature that is coming in Windows Server. This information will help you to get ready for Windows Server 2016 Technical Preview and give you an opportunity to develop and design a path to introduce this powerful technology into your environment and take full advantage of what is to come. This book was written at a time when the product was still evolving, and it should be noted that things might change or not appear in the final version of Windows Server 2016 when it is released. All guidance in these chapters is meant to be tried and evaluated in a test setting; you should not implement this in a production environment.

This book assumes that you are familiar with key Windows Server concepts (i.e., Hyper-V, networking, and storage) as well as cloud technologies such as Microsoft Azure. In this ebook, we cover a variety of concepts related to the technology and present scenarios with a customer focus, but it is not intended as a how-to guide or design manual. You should use other sources including the online Microsoft resources to stay up to date with the latest developments on the roles and features of Windows Server 2016 Technical Preview. The online resources will also contain the latest how-to procedures and information about designing a Windows Server 2016 infrastructure for your business.

Acknowledgments

We'd like to thank all the contributors who made this book possible:

- David Holladay
- Mitch Tulloch
- Ned Pyle
- Claus Joergensen
- Matt Garson
- John Marlin
- Robert Mitchell
- Deepak Srivastava
- Shababir Ahmed

- Ramnish Singh
- Ritesh Modi
- Jason M. Anderson
- Schumann Ge
- Yuri Diogenes
- David Branscome
- Shabbir Ahmed
- Ramnish Singh
- Andrew Mason
- The staff at Microsoft Press, who makes these titles possible!

Finally, to anyone we haven't directly mentioned, for all the help that has been provided, thank you!

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/WSpreview/errata>

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>.

The software-defined datacenter

Whether you're at a mid-sized business, a large enterprise, or a cloud service provider, you might think you've extracted as much value as you can from implementing virtualization in your datacenter. After all, virtualization saves your organization money by consolidating your server workloads and letting you retire obsolete hardware. You've also lowered your costs even further by moving some of your workloads from your on-premises infrastructure into a public or service provider–hosted cloud. But, you've also discovered that even though virtualization and cloud computing can address some IT challenges, they're no panacea—they also bring new problems. For instance, instead of dealing with server sprawl in your datacenter, you're now faced with virtual machine (VM) sprawl; in fact, it's worse because it's easier to spin-up a new VM than it is to procure and provision new server hardware in your environment. Instead of managing only one infrastructure in your datacenter, you now need to deal with managing resources in the cloud, as well. If you're not careful, you might end up with two sets of administrative tools that need twice as much time and staff to manage.

With this type of sprawl within the compute resource pool, we begin to have significant problems within the other pools of networking and storage. In essence, how do we keep up with what becomes cloud scale very quickly after an enterprise implements virtualization.

This is where a business of any size needs to begin thinking in terms of the *software-defined datacenter* (SDDC). As previously mentioned, the compute elements are pretty well covered and have been for nearly a decade or more. Now, however, we need to take new approaches to networking and storage and bring agile concepts to all three of the major resource pools: compute, storage and networking. This will ultimately make it possible for your environments to become true clouds and service the needs of the business on demand.

With this in mind, we can start on the path toward agile IT—toward building an IT infrastructure that is easy to grow and evolve as your business changes. Certainly if there's anything that characterizes the basic nature of doing business today, it's change.

Delivering agile IT

Most businesses think of SDDCs as hybrid clouds. This is because we are connecting our on-premises ecosystem to our cloud ecosystem. So, first let us define what exactly a hybrid cloud is. The National Institute of Standards and Technology, in a document titled "The NIST Definition of Cloud Computing" (<http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>), defines a hybrid cloud as the following:

Hybrid cloud—The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

But how exactly does this work? Suppose that your organization currently has a traditional datacenter with file servers, web servers, database servers, and so on, and you've also added on some cloud services from a hoster or public cloud provider so that you can run some of your applications and workloads in the cloud. Does that mean you already have a hybrid cloud?

No, it does not. The problem is that you're operating under two different paradigms: the old one of a traditional datacenter, which has very limited agility, and the new one of cloud computing, which has high agility and is managed very differently. Traditional datacenters are generally inflexible for a number of reasons and can scale up only slowly when the need arises. The cloud, on the other hand, is flexible and can scale easily on demand. Integrating these two different paradigms is like mixing oil and water: you can mix them if you try hard enough, but as time goes on, they will separate from each other.

The key to a hybrid cloud is to manage your resources as if the cloud were not just "out there," but also *in* your datacenter. When this is the case, you can easily move your compute, networking, and storage resources from on-premises to the cloud, and vice versa. This makes it possible for you to scale up or down quickly to address the changing needs of your business as the marketplace evolves. You can take advantage of solutions such as cloud backup when it makes sense for your business to do so. You also can use the same set of tools to manage and run your applications and services on-premises and in the cloud. This can make work a lot easier for everyone from your developers to your administrators, freeing up time so that they can concentrate their energy on important tasks to keep your business growing.

Understanding the Microsoft Cloud Platform

The Microsoft Cloud Platform can make all this possible for your business. Only Microsoft offers a consistent platform with which you can choose to run workloads where it makes the most sense for your business: in your datacenter, in a service provider's datacenter, or in Microsoft Azure. Figure 1-1 illustrates the breadth and benefits of the Microsoft Cloud Platform.




Figure 1-1: The Microsoft Cloud Platform

By using this platform, you can deliver and manage your IT services, both on-premises and in the cloud, in a unified way across a wide range of different device platforms. With the Microsoft Cloud Platform, you're able to do the following:

- Empower enterprise mobility
- Create the Internet of Your Things
- Foster application innovation
- Unlock insights on any data
- Transform the datacenter

The following sections briefly examine each of these benefits and how the Microsoft Cloud Platform delivers them.

Facilitating enterprise mobility

Making enterprise mobility possible is about the different kinds of devices that are used in the business environment today. It's about helping IT manage the needs of end users in their environment and about the governance and support of policies for devices, data, and users as those devices come into the enterprise. The Microsoft Cloud Platform facilitates enterprise mobility in a consistent platform through one pane of glass: Microsoft Intune.

More info To learn more about Microsoft Intune, go to <https://www.microsoft.com/en-ie/server-cloud/products/microsoft-intune/overview.aspx>.

Creating the Internet of Your Things

Creating the Internet of Your Things means that the Internet of Things (IoT)—what Forrester Research has equated with the concept of a connected world—starts with *your* things. In other words, it's not about ripping out and replacing technology in your enterprise; it's about taking advantage of what you have and using it in new ways, adding to your current systems, and innovating and optimizing so that everything works better together. It's about connecting existing devices and tapping into existing data. And, it's about getting away from spending all your time simply running your business and instead capturing time to think about how to make it thrive. The Microsoft Cloud Platform gives you the ability to do this securely through the Azure IoT Suite of services.

More info To read more about the Azure IoT Suite, go to <https://www.microsoft.com/en/server-cloud/products/sql-server/>.

Providing the right environment for application innovation

Fostering application innovation is about utilizing the power of the .NET environment. .NET is Microsoft's development environment that businesses can use to code applications once and then deploy them to any device. Businesses need modern capabilities to develop apps quickly through repeatable iterations using hybrid IT and cloud services. They also need apps that run smoothly on new form factors and incorporate built-in technologies such as business intelligence (BI) and social networking. The Microsoft Cloud Platform provides a comprehensive suite of software development tools and other technologies for building powerful, high-performance applications, including support for team-based design, development, and deployment, through Microsoft Visual Studio.

More info To learn more about Visual Studio, go to <http://msdn.microsoft.com/vstudio/>.

Unlocking insights on any data

Unlocking insights on any data is about helping IT bring together structured, relational data inside the enterprise with unstructured, nonrelational data from the web. This includes both Big Data from the web and small data in terms of enterprise resource planning (ERP), customer relationship management (CRM), and other apps within the firewalls—in other words, all data. It's about bringing together all of these different data sets to unlock insights for end users. The Microsoft Cloud Platform delivers the foundation for doing this in two ways: through Microsoft SQL Server and with Microsoft Power BI for Office 365. SQL Server gives businesses the tools to build mission-critical applications and Big Data solutions using high-performance, in-memory technology across online transactional processing (OLTP), data warehousing, BI, and analytics workloads, without having to buy expensive add-ons or high-end appliances. Power BI is a cloud service with which you can share, collaborate, and access your Microsoft Excel reports anywhere on any device.

More info To read more about SQL Server, go to <http://www.microsoft.com/server-cloud/products/sql-server/>. For more information on Power BI for Office 365, go to <http://www.microsoft.com/powerbi/default.aspx>.

Transforming the datacenter

Transforming the datacenter is about addressing the high levels of complexity involved in operating most datacenters today. Most of the energies of IT go into managing infrastructure, day after day. After this, the focus is on trying to consolidate servers and implementing virtualization to save money. What businesses actually need to do, though, is transform their datacenters by bringing in the cloud and by utilizing automation more effectively. Hybrid cloud makes it possible for IT to not simply think

about managing servers, but managing the pooled resources of compute, networking, and storage residing both in the datacenter and in the cloud and delivering these resources as shared services. The Microsoft Cloud Platform gives you the means to implement and automate hybrid cloud solutions through the combination of Windows Server, Microsoft System Center, and Azure.

More info To learn more about Windows Server, System Center, and Azure, go to <http://www.microsoft.com/server-cloud/> and then click Products.

How to transform your datacenter

The last of the five pillars—how you can transform your datacenter—is the focus of the remainder of this chapter. That’s because if you want to make your business more agile, you need to begin by focusing on your on-premises IT infrastructure. For most enterprises, this means looking carefully at the architecture and operations of the datacenter. In other words, you need to build an SDDC.

The term “software-defined” means that your IT infrastructure is decoupled from its underlying hardware so that you can manage and control it through policy. Virtualization is the key for doing this, and the Hyper-V virtualization platform of Windows Server provides businesses with the foundation for software-defined compute capability. Features such as Live Migration and Hyper-V Replica make VM mobility a reality and make it possible to decouple virtualized server workloads from the underlying physical server system fabric on which they are hosted.

But the Windows Server platform also provides more. Network Virtualization, a feature first introduced in Windows Server 2012, provides the underlying foundation for software-defined networking (SDN) capability with which you can create multitenant clouds on top of your underlying physical networking infrastructure. Storage Spaces, a technology also introduced in Windows Server 2012, provides the underlying foundation for software-defined storage (SDS) capability, making it possible for you to virtualize storage by grouping industry-standard drives into storage pools and then creating virtual drives, called storage spaces, from the available capacity in the storage pools. Each of these software-defined capabilities were later enhanced in Windows Server 2012 R2, and now with the impending release of Windows Server 2016, these capabilities have been expanded and improved to make them more powerful and flexible than ever.

However, Windows Server provides only the foundation for implementing software-defined compute, networking, and storage capabilities in your datacenter. To realize the full benefit from these capabilities by automating them, you need System Center—in particular System Center Virtual Machine Manager—as well as the Azure Pack, a collection of Azure technologies with which you can bring the functionality and manageability of Azure into your datacenter. System Center and the Azure Pack together with Windows Server make the SDDC possible. Finally, to extend your SDDC to the cloud, you need Azure—Microsoft’s global, enterprise-grade cloud platform that offers compute, storage, data, networking, and app services—as well as the innovative Azure capabilities that are built in to Windows Server 2016. These capabilities are coming in the next version of System Center, which will be released in conjunction with Windows Server 2016.

Windows Server 2016, together with Azure and the upcoming versions of System Center and Azure Pack, are the building blocks you can use to build the SDDC that makes agile IT possible through the hybrid cloud; they form the foundation of the Microsoft Cloud Platform.

The next section of this chapter focuses on the exciting new features of Windows Server 2016—many of which will work in conjunction with System Center and Azure—that can transform your datacenter by making it possible for you to do the following:

- Build a software-defined foundation
- Automate and secure your infrastructure
- Extend your infrastructure to the cloud on demand

Building a software-defined foundation

Figure 1-2 illustrates that there are three basic parts in the architecture of an SDDC. You can think of these different parts as “planes” because each spans the entire infrastructure while complementing one another.




Figure 1-2 Abstraction layers of the software-defined datacenter

From the bottom up, the three planes of the SDDC are as follows:

- **Data plane** This plane represents an abstraction of the underlying compute, networking, and storage hardware of the physical infrastructure. The data plane is where network traffic actually flows on and through hardware, such as switches, routers, server systems, storage devices, and so on. It's important to understand that in the software-defined paradigm, these hardware devices are not the data plane itself; instead, the data plane consists of various abstractions of the underlying hardware. For example, a VM is an example of a compute resource on the data plane, whereas a virtual network is an example of a network resource.
- **Control plane** This plane represents an abstraction of the portion of the infrastructure that controls how the network traffic flows on the data plane. You can think of the control plane as the software and protocols that coordinate the resources on the data plane and provide decision logic to ensure that the datacenter functions as a distributed system as intended. The control plane is therefore a higher level of abstraction than the data plane. An example of a control plane element is the Hyper-V Network Virtualization (HNV) functionality introduced in Windows Server 2012 and built in to succeeding releases of the Windows Server platform.
- **Management plane** This plane represents an abstraction of the portion of the infrastructure with which you create, deploy, manage, monitor, and maintain policy for elements of the control plane. The management plane is thus the highest level of abstraction for the software-defined datacenter. An example of a management plane element is System Center Virtual Machine Manager, which you can use to configure and manage the VM networks, virtual subnets, logical networks, and other abstractions needed to implement HNV in a datacenter.

The key element of an SDDC architecture is that all three of these planes—the management, control, and data planes—are separated from one another and can be controlled through software. As such, a business can manage its information systems resources holistically. This also makes the infrastructure flexible, resilient, and agile. Windows Server along with System Center and Azure make it possible for organizations to set up SDDCs. This transformation began with Windows Server 2012, and the technologies were enhanced in Windows Server 2012 R2. Windows Server 2016, with its tighter integration with Azure, and the next version of System Center make implementing the software-defined paradigm in your datacenter easier than ever before.

What makes it possible

How will Windows Server 2016—combined with Azure and the upcoming release of System Center—make the SDDC architecture possible? This section introduces the new features and enhancements in Windows Server 2016 and directs you to where you can find out more about these improvements.

To describe the improvements to the Windows Server platform (which is where Hyper-V virtualization resides) and the virtualized workloads (VMs) that run on Hyper-V hosts, this section focuses on the data plane of the SDDC. This gives us an opportunity to examine the new features and enhancements of Windows Server 2016 from three different perspectives:

- **Compute improvements** Compute improvements include new features and enhancements that organizations can use to confidently virtualize their enterprise workloads. Organizations need enterprise-grade reliability and flexibility and a platform that runs traditional, distributed, and cloud applications. To meet these needs, Windows Server 2016 delivers best-in-class scale, performance, and resilience for enterprise workloads; the ability to deploy and manage Linux as a first-class element of your infrastructure; and a frictionless fabric that you can upgrade without downtime.
- **Networking improvements** Networking improvements include new features and enhancements that make flexible workload placement and mobility possible. Organizations need flexibility, reliability, high levels of performance, and a focus on applications and workloads. To meet these needs, Windows Server 2016 delivers enhancements in the reliability, performance, and interoperability of virtual networking; improved support for centralized configuration and management across virtual and physical networks; new virtualized network functions for transforming the network cloud; and seamless datacenter extensions for flexible workload placement and mobility.
- **Storage improvements** Storage improvements include new features and enhancements that reduce enterprise storage costs. Organizations need increased storage efficiency, improved protection of key data and workloads, and the ability to easily grow storage capacity—all while keeping storage costs under control. To meet these needs, Windows Server 2016 delivers the ability to deploy a cost-effective, cloud-scale, SDS platform; centralized deployment and management for more resource-efficient on-premises storage; a provision for delivering business continuity for data and workloads; and the ability to deploy a hybrid cloud storage solution.

Compute improvements

Some of the key compute improvements in Windows Server 2016 that facilitate the SDDC include the following:

- **Hot add and remove for network adapters and memory** With this new feature, you can add or remove a network adapter and adjust the amount of memory assigned while the VM is running, without any interruption. The memory adjustment capability even works when you have Dynamic Memory turned on for a Hyper-V host. For more information about this feature, go to http://technet.microsoft.com/library/dn765471.aspx#BKMK_hot.

- **Rolling cluster upgrades** Using this feature, you can add a node running Windows Server 2016 Technical Preview 2 to a Hyper-V cluster with nodes running Windows Server 2012 R2. For more information about this new feature, go to http://technet.microsoft.com/library/dn765471.aspx#BKMK_HyperVRollingUpgrades. You might want to review the topic “Cluster Operating System Rolling Upgrade in Windows Server Technical Preview” at <http://technet.microsoft.com/library/dn850430.aspx>.
- **True VM mobility** Hyper-V in Windows Server 2016 introduces support for down-level live migration. Chapter 2 includes a walkthrough of this new feature.
- **Virtual Machine Groups** This new feature is designed to make the management of multiple VMs easier. Chapter 2 describes this in more detail.
- **VM configuration version** The VM upgrade process has changed in Windows Server 2016 (see Chapter 2 for more information).
- **New configuration file format** Hyper-V in Windows Server 2016 now uses a new configuration file format that reduces the chances of data corruption should the VM storage for your Hyper-V host fail. You'll learn more about this new feature in Chapter 2.
- **Production checkpoints** Windows Server 2016 supports taking checkpoints for production VMs running Microsoft Windows. This new capability uses the Volume Snapshot Service of the guest operating system, as described in Chapter 2.

Networking improvements

Some of the key networking improvements coming in Windows Server 2016 include the following:

- **Network Controller** This new feature provides a centralized, programmable point of automation to manage, configure, monitor, and troubleshoot virtual and physical network infrastructure in your datacenter. For more information, go to <http://technet.microsoft.com/library/dn859239.aspx>.
- **Distributed multitenant firewall** This new feature protects the network layer of virtual networks. For learn more, go to <http://technet.microsoft.com/library/dn859239.aspx>.
- **Support for GRE tunneling** This new feature turns on a Generic Routing Encapsulation (GRE) tunnel capability for the Windows Server Gateway. For read more, go to <http://technet.microsoft.com/library/dn765485.aspx>.
- **Service chaining** This new feature provides the tenant administrator with the means to specify multiple virtual appliances in a chain to be grouped together so that the selected network traffic passes through each of these appliances in the order specified. For more information, go to <http://technet.microsoft.com/library/dn859239.aspx>.
- **Multitenant edge gateway** This new feature provides a multitenant gateway solution that gives tenants the tools to access and manage their resources over site-to-site VPN connections from remote sites. It also allows network traffic flow between virtual resources in the cloud and tenant physical networks. For more information, go to <http://technet.microsoft.com/library/dn859239.aspx>.

In addition, some networking improvements involving Azure are relevant to building an SDDC:

- **Software load balancer based on Azure** This feature is a Layer-4 load balancer that represents a version of the Azure offering and has been deployed at scale in the Azure environment. For more information, go to <http://technet.microsoft.com/library/dn890699.aspx>.

- **Azure ExpressRoute** This feature of Azure is used to create private connections between Azure datacenters and on-premises infrastructure or in a colocation environment. This is covered in more detail later in this chapter.

Storage improvements

Some of the key storage improvements in Windows Server 2016 that can help you implement the SDDC include the following:

- **Storage replication for any volume** Storage Replica is a new feature that provides storage-agnostic, block-level, synchronous replication between servers for disaster recovery, and allows stretching of a failover cluster for high availability. For more information on this new feature, go to http://technet.microsoft.com/library/dn765474.aspx#BKMK_SR. Chapter 3 includes some technical information about Storage Replica.
- **Shared Nothing Storage** With this new feature, you can build highly available storage systems using only local storage instead of needing to use drives that are physically connected to all storage nodes. You can even do this with SATA drives, something that was not previously possible with Windows Server. Chapter 3 presents more about Shared Nothing Storage.
- **Storage Quality of Service** This new feature gives you the ability to create storage QoS policies on a Scale-Out File Server and assign them to one or more virtual drives on Hyper-V VMs. For more information about this new feature, go to http://technet.microsoft.com/library/596f28ec-e154-4c2e-9e82-7e42afe0e9fa#BKMK_QoS. You might also want to review the post titled “Storage Quality of Service Guide Released for Windows Server Technical Preview” on Jose Barreto’s blog at <http://blogs.technet.com/b/josebda/archive/2014/10/24/storage-quality-of-service-guide-released-for-windows-server-technical-preview.aspx> and view the Windows Server Technical Preview Storage QoS Guide at <http://go.microsoft.com/fwlink/?LinkId=517877>.
- **Cloud Witness using Azure** Failover Clustering in Windows Server 2016 supports using a witness that you create in the cloud by using Azure. Chapter 3 includes a walkthrough of this new feature.
- **Shared VHDX improvements** Windows Server 2016 includes some improvements to the Shared VHDX feature previously introduced in Windows Server 2012 R2. Chapter 3 describes these improvements more fully.
- **Improved cluster logs** The Cluster Diagnostic Log has been improved in Windows Server 2016 to make it easier to use the log for troubleshooting. You can learn more about these improvements in Chapter 3.

In addition, some storage improvements involving Azure are relevant to building a SDDC:

- **Azure Backup** This cloud-based backup solution provides a reliable, inexpensive, and scalable solution with zero capital investment and minimal operational expense. For more information, go to <http://azure.microsoft.com/services/backup/>.
- **Azure Site Recovery** This feature protects important applications by coordinating the replication and recovery of physical or virtual machines. Recovery is described in more detail later in this chapter.
- **Azure StorSimple** This integrated storage solution manages storage tasks between on-premises devices and Azure cloud storage. This feature is covered in more detail later in this chapter.

Automating and securing your infrastructure

Building a software-defined foundation isn't the only thing you'll need to do to transform your datacenter. Automation and management are also important for achieving operational excellence for your organization's IT processes and services. In Windows Server 2016, the following are some of the key improvements in this area:

- **Shielded VM** This feature ensures that the data and state of VMs are protected at all times using proven security technologies. You'll find more information about this feature in Chapter 5, "Security."
- **Control Flow Guard** This feature ensures that software written to take advantage of Control Flow is automatically protected against attacks such as buffer overflows. This feature is covered in more detail in Chapter 5.
- **Device Guard (Code Integrity)** This feature ensures that only software authorized by you can run on a machine, helping protect against unauthorized applications. This feature is covered in more detail in Chapter 5.
- **Windows Defender** This feature, turned on by default, helps protect machines from malware from the moment you install the Windows operating system. The feature is covered in more detail in Chapter 5.
- **Enhanced Audit** This feature adds two new categories to provide deeper levels of auditing to find suspect behavior. This is covered in more detail in Chapter 5.
- **Just Enough Administration/Just In Time** This feature gives you the ability to provide only the privilege that is required, when it is required. This feature is covered in more detail in Chapter 5.
- **PowerShell V5** This feature provides enhanced scripting capabilities for configuration, management and deployment in an SDDC. This feature is covered in more detail in Chapter 7, "Systems management."

In addition, Microsoft System Center is a key tool that helps an organization to realize the full benefits of the Microsoft Cloud Platform. System Center delivers unified management across on-premises, service provider, and Azure environments. Through this unified management, you can do the following:

- Deliver higher levels of infrastructure and application resiliency
- Reduce complexity by simplifying how you provision, manage, and operate your infrastructure
- Drive efficiencies in your IT operations and processes through integrated automation

And, by automating your IT operations and implementing standardized best practices, you can do the following with System Center:

- Simplify cloud management with consistent processes
- Improve operations through infrastructure, workload, and application monitoring
- Streamline provisioning by cloud and workload deployment and configuration

There can be challenges, however, in achieving the goal of automating and securing your infrastructure. The biggest challenge is that business and IT live in very different worlds.

IT strives above all for stability. That's because the key business expectation of IT is that applications and services just work. To achieve such stability, IT relies heavily on control. To that end, many of the policies and procedures that IT puts into place are designed to control the environment.

Business, on the other hand, strives above all else for growth. Because the marketplace constantly changes, businesses need to be flexible and agile so that they can quickly evolve in response to changing economic conditions. This puts tremendous stress on traditional IT processes, which by design change slowly to maintain stability. Business says, "We need X to grow, and we need it now!" IT responds, "Wait! We'll need to test this and introduce it slowly to make sure it doesn't upset things." As a result, business is frustrated and tries an end-run around IT by implementing what's known as *shadow IT*—IT solutions implemented without the blessing or knowledge of the IT department.

Of course, shadow IT has been around for as long as computers have been affordable and readily available for businesses. Copying company data onto removable storage to take home is an example of end users doing an end-run around the controls IT puts in place to safeguard sensitive business information. The endless proliferation of Microsoft SharePoint sites can also be considered a form of shadow IT because such sites are usually self-provisioned without any oversight from the IT department. Other examples of shadow IT range from setting up unauthorized Wi-Fi access points to deploying entire Active Directory domains of PCs with their own DHCP servers.

And then there's the cloud.

One can argue that the main driver behind business units surreptitiously signing up for cloud services with a public or hosted-cloud provider is the backward-looking, overly-restrictive resistance to change evidenced by IT departments in large organizations. Because of such inertia, and because employees are driven to perform in order to be rewarded (or at least not penalized), it's becoming commonplace for those whose workflow can benefit from using cloud-based services to secretly implement them in the office.

The problem with this scenario is that parallel IT infrastructures that operate in the shadows are not parallel at all. They touch the company's own infrastructure at various points; where they touch can represent vectors for new forms of attack on the company's business assets. The problem is magnified yet further by how cloud services are becoming cheaper, more powerful, and easier to sign up for and use.

The solution, of course, isn't to avoid using the cloud but to integrate cloud service delivery models such as Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) into the existing IT infrastructure of your organization. Cloud service delivery provides the much needed interface between IT and business, making it possible for each to operate in its own style but in a compatible way. Standardized IaaS and PaaS services are a key factor in eliminating much of the custom IT deployments that previously cost so much time and money in traditional IT. And the results of implementing integrated cloud service delivery are greater agility for business, continued control for IT, and flexibility to both business and IT in terms of which types of tools, technologies, and cloud services are used.

Extending to the cloud on demand

So far, this chapter has described how changes in your on-premises infrastructure and a software-defined approach can make your IT infrastructure (and therefore your business) more agile. It has also described how Azure Stack will make this hybrid world an easy and manageable reality. Now, let's look briefly at how hybrid capabilities can help solve your existing problems and make your job easier.

At the beginning of this chapter, we said that a hybrid cloud is the solution to the dilemma of needing to manage two separate infrastructures: the one in your datacenter and the one in the cloud. That's because the hybrid approach makes it possible for you to integrate your on-premises infrastructure

with compute, networking, and storage resources in the cloud. We also emphasized that a hybrid cloud is the path toward agile IT—toward building infrastructure that can grow and evolve as your business changes. We said that the key to implementing a hybrid cloud solution is being able to manage your IT resources as if the cloud were not only “out there,” but also right inside your datacenter. Because if the cloud is also within your datacenter, you will be able to quickly and easily move compute, networking, and storage resources from on-premises into the cloud and from the cloud back into your datacenter.

The Microsoft cloud platform already includes numerous capabilities that make extending your datacenter to the cloud on-demand possible. Here are some of the key capabilities:

- **VM mobility** By using Windows Server Hyper-V features such as Live Migration and Storage Migration, and by using System Center products like Virtual Machine Manager and App Controller, you can easily move VMs between private, hosted, and Azure clouds. And, as mentioned previously, Hyper-V in Windows Server 2016 also supports down-level live migration. For a good overview of the current support for VM mobility in the Windows Server Hyper-V platform, download the Hyper-V Component Architecture Poster from the Microsoft Download Center page at <https://www.microsoft.com/download/en/details.aspx?id=29189> and see also the TechNet article titled, “Poster Companion Reference: Hyper-V Virtual Machine Mobility,” which is available at <https://technet.microsoft.com/library/dn641214.aspx>.
- **Azure Hybrid Use Benefit** This benefit is offered when you have enhanced your purchase of Microsoft Windows Server with Software Assurance (SA). With Azure Hybrid Use Benefit, you can use the non-Windows rate for Windows Server instances in Azure. If you choose either Windows Server Datacenter or Windows Server Standard, you can get either two instances of 1 to 8 cores or one instance of up to 16 cores; however, with Standard you cannot use the license for Windows Server on-premises at the same time. With Datacenter, you can run a server on-premises and in the cloud at the same time. To read more about Azure Hybrid Use Benefit, go to <https://azure.microsoft.com/en-us/overview/azure-for-microsoft-software/faq/>.
- **Azure StorSimple** StorSimple is a hybrid online cloud storage service for enterprises that is designed to reduce costs and improve data protection. StorSimple provides bottomless cloud storage connected to Azure. Installing the StorSimple hybrid storage array in your datacenter applications in physical or virtual servers can make a local network connection that automatically tiers to Azure storage with inline deduplication and compression. This offers a range of usage scenarios including file shares, SQL servers, VMs, SharePoint, and data archiving. For information about StorSimple, go to <http://azure.microsoft.com/services/storsimple/> and also see the overview on the Microsoft Cloud Platform site at <http://www.microsoft.com/server-cloud/products/storsimple/>.
- **Azure ExpressRoute** You can extend your datacenter into Azure in either of two ways. First, you can simply have your encrypted traffic traverse across the public Internet, in which case the performance level will depend greatly on the available bandwidth allowances over which you have little control. The second way is to use a dedicated network for which you can fully manage and audit the security and throughput from one end of the connection to the other. ExpressRoute does just that by giving you a fast and reliable connection to Azure, which makes it suitable for scenarios like periodic data migration, replication for business continuity, disaster recovery, and other high-availability strategies. For more information on ExpressRoute, go to <http://azure.microsoft.com/services/expressroute/>.
- **Azure Active Directory** Active Directory is the foundation for how hybrid identity operates in the Microsoft Cloud Platform. For identity and access, the key is the ability to maintain a single identity across multiple clouds. Continuous services and connected devices present a real challenge, with users expecting more and more from IT in terms of simple and fast access to resources and data. Microsoft offers several options in this area, including the advances in identity management in both Windows Server Active Directory and Azure Active Directory, a cloud

solution that provides capabilities for managing users and groups to secure access to your on-premises and cloud applications, ranging from Office 365 to non-Microsoft Software-as-a-Service (SaaS) applications. Cloud-based identity that integrates with your existing Active Directory solution provides tremendous flexibility in building single sign-on capabilities across your cloud deployments. For more information on Azure Active Directory, go to <http://azure.microsoft.com/services/active-directory/>.

- **Azure Site Recovery** Azure Site Recovery protects your applications by coordinating the replication and recovery of physical or virtual machines. With Azure Site Recovery, you can easily replicate a VM within a datacenter to a hosting service provider or to Azure. Site Recovery provides three key capabilities for your environment:
 - **Simple, automated protection** You can protect your environment by automating the replication of the VMs based on policies that you set and control. Site Recovery coordinates and manages the ongoing replication of data by integrating with existing technologies such as Hyper-V Replica, System Center, and SQL Server AlwaysOn.
 - **Continuous health monitoring** Site Recovery monitors the state of Virtual Machine Manager clouds continuously and remotely from Azure. When replicating between two sites you control, only the Virtual Machine Manager servers communicate directly with Azure—your VM's data and replication remains on your networks. All communication with Azure is encrypted.
 - **Orchestrated recovery** The service helps automate the orderly recovery of services in the event of a site outage at the primary datacenter. You can bring up VMs in an orchestrated fashion to help restore service quickly, even for complex, multitier workloads. Recovery plans are simple to create through the Azure Management Portal, where they are stored. The plans can be as simple or as advanced as your business requirements demand, including custom Windows PowerShell scripts and pauses for manual interventions. You can also customize networks by mapping virtual networks between the primary and recovery sites. You can test these plans whenever you like without disrupting the services at your primary location.

More info To learn more about Azure Site Recovery, go to <https://azure.microsoft.com/services/log-analytics/>, which refers to Ops insights versus log-analytics.

In addition to these features, the Microsoft Cloud Platform also provides various hybrid management capabilities to make it easy for you to extend the management of your infrastructure into the cloud. These hybrid management capabilities function in three areas: operations, automation, and insight. For operations, capabilities include intelligent workload and platform monitoring across hybrid environments. Automation is implemented through a rich workflow by delivering a single automation solution across clouds. Cloud-based operations management with an intelligence service lets you gain deeper insights into your environment.

Examples of some of the hybrid management capabilities already available through Azure include the following:

- **Azure Operational Insights** This is a SaaS-based service, intended for IT operations teams, that uses Azure HDInsight to collect machine data across environments and turn it into real-time operational intelligence to assist you in making better-informed business decisions. Operational Insights includes ready-made intelligence packs and integrated search capability. For more information on Azure Operational Insights, go to <https://azure.microsoft.com/services/log-analytics/>.

- **Azure Automation** With this service, you can create, deploy, monitor, and maintain resources in an Azure environment by using a highly scalable and reliable workflow engine. You can use this Azure-based service to orchestrate time-consuming and frequently repeated tasks, decreasing operational expense for cloud operations. For more information about Automation, go to <http://azure.microsoft.com/services/automation/>.
- **Visual Studio Online** DevOps integration is supported through connections to Visual Studio Online for alignment of both release management and application performance monitoring. For more information on Visual Studio Team Services, go to <https://azure.microsoft.com/services/visual-studio-team-services>.

Compute

In this chapter, we explore all of the new features related to Compute in Windows Server 2016 Technical Preview. Specifically, we focus on Failover Cluster, Hyper-V, and Multipoint Services.

Because this is still a technical preview, features may change or be removed before the product becomes generally available. Throughout this chapter we will refer to Windows Server 2016 Technical Preview with respect to its current release, that being Version 4.

Failover Clustering improvements

By John Marlin and Colin Robinson

This section describes a key improvements Microsoft has made to Failover Clustering in Windows Server 2016 Technical Preview. These improvements include the following:

- Cloud witness using Microsoft Azure
- Shared VHDX improvements
- Improved cluster logs
- Active memory dumps
- Network name diagnostics
- Cluster Operating System Rolling Upgrade

In addition, Colin Robinson walks us through a detailed demonstration of performing a rolling upgrade on a failover cluster.

Creating a cloud witness by using Azure

Beginning with Windows Server 2008, each version of Failover Clustering has introduced a new quorum type to complement what already exists. That hasn't changed. Windows Server 2016

Technical Preview introduces the cloud witness quorum type, a witness that you can create in the cloud by using Azure.

This quorum type takes advantage of the Azure public cloud as the arbitration (witness) point for the cluster. You can achieve this configuration without the need for an extra site and you will utilize it mostly in multisite clusters. It provides a quorum option for the following situations:

- Multisite clusters that do not have a third site on which to place a file-share witness
- Clusters using nonshared storage
- Guest clusters hosted in Azure
- Guest clusters hosted in private clouds
- Clusters using direct-attached storage (DAS)

The cloud witness acts the same as a file-share witness, using the same basic logic in that it does not contain a copy of the cluster database and will act as a deciding vote to prevent *split brains* (multiple nodes running in the same cluster that cannot communicate with one another).

To configure a cloud witness, you first must have an Azure subscription. Here are the steps you need to take to get one:

1. Sign in to the Azure management portal and create a storage account for this witness, as shown in Figure 2-1.




Figure 2-1: Creating a storage account in Azure

2. When the storage account is created, highlight it in the portal and click Manage Access Keys. Copy the primary access key for later use.
3. In the Failover Cluster Manager console, configure the quorum for the cloud witness. First, right-click the name of the cluster and then, on the shortcut menu that opens, click More Actions, and then select Configure Cluster Quorum Settings, as demonstrated in Figure 2-2.




Figure 2-2: Configuring cluster quorum settings

4. The Configure Cluster Quorum Wizard opens. On the Select Quorum Configuration Option page, click Select The Quorum Witness, as depicted in Figure 2-3.




Figure 2-3: Selecting the Quorum Type

5. On the Select Quorum Witness page, click Configure A Cloud Witness, as illustrated in Figure 2-4.




Figure 2-4: Configuring a cloud witness

6. On the Configure Cloud Witness page, type the storage account name you created in the management portal, your Azure primary storage account key, and the Azure service endpoint, as shown in Figure 2-5.




Figure 2-5: Entering storage account information

Note You can do the same in Windows PowerShell by using the following command:

```
Set-ClusterQuorum -CloudWitness -AccountName MyWitness -AccessKey <your storage key> -Endpoint core.windows.net
```

There are two key prerequisites for using the cloud witness:

- You must have a valid Azure subscription.
- All nodes must have Internet access and be able to access Azure.

Additionally, as with a file-share witness, you can use the same Azure account or container for multiple clusters.

Shared VHDX improvements

Since Windows Server 2008 Hyper-V, you have had the ability to create guest clusters as virtual machines (VMs). However, to have any sort of shared storage, you were required to use iSCSI. Windows Server 2012 introduced Virtual Fibre Channel support for VMs as a second option for shared storage.

However, from the perspective of a service provider, Virtual Fibre Channel is not always a viable option. Virtual Fibre Channel opens and provides the customer with access to the physical storage infrastructure in the same way that physical iSCSI does. However, if a service provider set up a VM and added iSCSI support for the customer-shared drives, the customers might be unhappy because they would be charged for an additional VM.

Because of these concerns, Microsoft introduced Shared VHDX in Windows Server 2012 R2 as an additional option. Shared VHDX gives guest clusters the shared storage they needed without access to storage infrastructures. This did add another option from a shared-drive perspective; however, it was not without limitations. In the latest Windows Server 2016 Technical Preview, improvements have been made to address some of these limitations.

Suppose that you have a Shared VHDX drive that is filling up, and you need to increase the size. In Windows Server 2012 R2, downtime was unavoidable because to increase the size the VMs would need to be powered off. That is not an ideal solution for a 24/7 business. In Windows Server 2016 Technical Preview, you can now expand the drive while it is online. (Note that you can only expand a Shared VHDX drive, you cannot shrink one.)

To expand the drive, perform the following steps:

1. Open Failover Cluster Manager, right-click a VM, and then select Settings.
2. Click the drive that you want to expand, as depicted in Figure 2-6.




Figure 2-6: Settings for Hard Drive

- Click Edit to start the Edit Virtual Hard Disk Wizard. The only available option is Expand, so the selection is already made for you, as illustrated in Figure 2-7.




Figure 2-7: Expanding a hard drive

- On the Configure Disk page, type the size you want the virtual hard drive to be, as shown in Figure 2-8.




Figure 2-8: Type the size for expansion

Note You can do the same by using the following Windows PowerShell cmdlet:

```
Resize-VHD -Path C:\ClusterStorage\Volume4\Shared-VHDX\Disk2.vhdx -SizeBytes 32212254720
```

- When you have completed the wizard, open the VM and expand the drive in Server Manager, as demonstrated in Figure 2-9.




Figure 2-9: Extending Volume in Server manager

With your VMs all running on your Hyper-V cluster, it is time to back up the machines. In Windows Server 2012 R2, you cannot back up the Shared VHDX attached to a VM from the host. Because it is shared, it is blocked from being backed up. However, in Windows Server 2016 Technical Preview, you can select it as a VHD to back up.

VMs that include a Shared VHDX can now also participate in Hyper-V Replica. In previous versions of Windows Server, VMs with a Shared VHDX were blocked from participating. With the improvements in Shared VHDX, you not only have the option to replicate the VMs, you also have the option to select any or all of the Shared VHDX drives, as presented in Figure 2-10.




Figure 2-10: Choosing Replica VHDs

Improved cluster logs

Getting as much critical data as possible in a timely fashion can help to quickly resolve Failover Clustering problems. Getting the right data at the right time can be critical for restoring services. With all the SLAs available, the right diagnostics is crucial to many businesses.

From a diagnostic standpoint, the first improvement in Windows Server 2016 Technical Preview is with the cluster diagnostics log. The cluster log has always been useful for identifying an error, what led to an error, what was happening at the time of the error, and so on. But using the cluster log to determine things such as the configuration of the cluster is a more involved process.

For example, if you want to know the resources in the cluster, as long as the cluster log shows the cluster service starting, the required information is there, but finding it requires searching through many lines to piece it all together, as demonstrated in the following example:

```
<Networks><vector len='4'>
<item><obj sig='NETW' id='1f509983-2478-4630-af8a-e13d2c486172' name='Cluster Network 2'/></item>
<item><obj sig='NETW' id='967df8d8-0f94-4d02-93d5-046fa1ce2369' name='Cluster Network 1'/></item>
<item><obj sig='NETW' id='99e6e621-0de5-4a1c-a468-a68057ee6278' name='Cluster Network 4'/></item>
<item><obj sig='NETW' id='a171b564-6b89-4c7d-91a5-f2dcbe450fbc' name='Cluster Network 3'/></item>
</vector>

<LIVE id='.Live' name='.Live'>
<NODE id='2' name='2012R2N1'>
<ITFC id='62edcfaa-fb25-4108-ac2b-236b3520af3f' name='2012R2N1 - WAN'/>
<ITFC id='884676e9-6df6-4a26-a423-c9b113a99056' name='2012R2N1 - iSCSI 2'/>
<ITFC id='65fa8c93-c242-45f0-88ab-9e26f0845c0b' name='2012R2N1 - Public'/>
<ITFC id='10962cca-3015-4380-8011-76e47ef575c4' name='2012R2N1 - iSCSI 1'/>
</NODE>
<NODE id='1' name='2012R2N2'>
<ITFC id='9e79c4c8-8d78-4d14-ab51-7a39d7191c4a' name='2012R2N2 - WAN'/>
<ITFC id='09569c99-262d-4f6b-95e4-f69185669f2b' name='2012R2N2 - iSCSI2'/>
<ITFC id='aa8f05e2-eeae-452c-9960-3905d0319fe1' name='2012R2N2 - Public'/>
<ITFC id='95d0074f-6ff6-4b5f-89a7-454fe2306332' name='2012R2N2 - iSCSI1'/>
</NODE>
</LIVE>
```

Getting just this little bit of information entails going through 1,000 or more lines in the logs and can be time consuming. Looking up other specific information about the configuration can take you back through these same sets of entries plus more.

You can get this type of information from other logs or the registry, but that requires reviewing information from multiple files. If you are not at the machine and someone else gathers logs for you, that person might not include all the logs. This can delay things further as you wait to get the proper information.

This complexity was an important consideration for Windows Server 2016 Technical Preview as it relates to Failover Clustering. Because of this, the cluster diagnostics log has been redesigned. When you generate a cluster log in Windows Server 2016 Technical Preview, it includes additional information that can be accessed quickly, broken down into various sections, as illustrated here:

```
[=== Cluster ===]
This section gives information about the version, time running, node name the log came from, etc

[=== Resources ===]
List of all resources (including the GUID) and the configurations/parameters of those resources

[=== Groups ===]
List of all groups (including the GUID) and the configurations/parameters of those groups, owner node, etc

[=== Resource Types ===]
List of all resource types (including the GUID) and the configurations/parameters of those resource types

[=== Nodes ===]
This section gives information about the nodes including the version, time running, node id, etc

[=== Networks ===]
This section gives information about the networks including the role, the network scheme, metric, if RSS capable, etc

[=== Network Interfaces ===]
This section gives information about the networks including the name, IP Address information, etc

[=== System ===]
All System Event Log entries with Failover Clustering as the source

[=== Microsoft-Windows-FailoverClustering/Operational logs ===]
All events from the Microsoft-Windows-FailoverClustering/Operational channel that gives you information about the forms of a cluster, node joins, group moves, etc

[=== Microsoft-Windows-ClusterAwareUpdating-Management/Admin logs ===]
All events from the Microsoft-Windows-ClusterAwareUpdating-Management/Admin channel that gives you information about Cluster Aware Updating
```

```
[=== Microsoft-Windows-ClusterAwareUpdating/Admin logs ===]
All events from the Microsoft-Windows-ClusterAwareUpdating/Admin channel that gives you information about
Cluster Aware Updating

[=== Microsoft-Windows-FailoverClustering/DiagnosticVerbose ===]
This is actually a new event channel that gives you the similar output as Debug Level 5 for a Cluster Log
without having to set it. You can use this information to get deeper into the calls and goings on with the
cluster and gives it in a more verbose output.

[=== Cluster Logs ===]
This is the output that you normally see in a cluster log.
```

Clearly, there is a lot to this log now. Using this one log can reduce the time you spend looking for information or trying to resolve an issue. Instead of having to review three or more files that might be loaded in three different applications with their own formats, you now need only review one file.

The other useful thing about the log is that if you generate it with no switches, everything in an event channel (for example, the System Event) is shown. So you can actually get any history pertaining to a particular problem. In cases for which you can reproduce an error condition and do not need all of the history, you can generate a log for the past five minutes (`TimeSpan=5`). Helpfully, the new cluster log uses this same five-minute time span for all of the event channels and gives you only those, so you don't need to deal with an unnecessarily large file.

Active memory dump

Another new feature as it relates to diagnostics is the ability to capture memory dumps. Imagine that you have a big Hyper-V cluster and each of the nodes has 512 GB of memory. If the node is having an issue and you create a memory dump that contains both user and kernel mode memory, that memory dump is going to be over 512 GB. Trying to work with a file that size can be a nightmare. First, you need to ensure that you have a drive with enough free space to hold a file of that size. You then need to spend hours zipping, uploading, and unzipping before you can even begin to open it. If the problem is with the host server itself and you are running VMs that use 500 GB of that memory, this is information that you don't need, because it does not pertain to the host.

Because of this, there is a new dump setting called Active Memory Dump. This setting captures only the memory that the host is actually using. If the host is actively using only 5 GB of memory, a 5-GB memory dump is what will be created. This smaller dump is much easier to parse than the 512-GB file in the previous scenario.

The Active Memory Dump option is in the same location as the normal dump settings, in the Startup And Recovery dialog box of the System Properties, as shown in Figure 2-11.




Figure 2-11: Memory dump settings

Note On the topic of memory dumps in general, Failover Clustering has been integrated with Live Dump to capture dumps when timeouts are reached. You then can analyze these dumps for root cause. This also integrates with Windows Error Reporting to allow for retention and inclusion of other logs.

Depending on how a resource times out and how Failover Clustering is configured, Live Dump can cause the machine to issue a stop error (i.e., blue screen) and cause the machine to stop responding in order to create a memory dump. Although the memory dump is useful in determining the cause of a problem, it does incur downtime while the machine reboots to create the dump. However, with the integration of Live Dump, a memory dump is created in the background while the machine itself continues to run, which does not affect production.

The focus of this feature is to collect enough logs/dumps for Microsoft support to more successfully troubleshoot various issues that customers might experience when using clusters. The goal is to gather enough information so that support can help solve the problem when customers call, instead of asking customers to reproduce the problem and then waiting for a time when it can be done.

Network name diagnostics

Windows Server 2016 Technical Preview features improvements related to diagnosing network name problems. At times, some of the events are confusing or not even present. For example, previously, problems updating DNS resulted in a generic error indicating that DNS could not be updated. But the error notification did not indicate why DNS could not be updated. Several things could prevent DNS from being updated:

- DNS does not accept dynamic updates.
- You are using a secure DNS server and the cluster does not have the proper rights.
- There are timeouts getting to the DNS server.

Troubleshooting an event such as this took time because you had to look at it in broad terms before you could focus on a specific area and narrow down the problem. In Windows Server 2016 Technical Preview, events are updated to include the specific error. So, if the error is due to one of the reasons mentioned in the list, you are notified of the error and can immediately focus on that one cause. This makes for quicker resolutions because you do not need to troubleshoot a problem that does not exist.

Also, additional checks have been added for the network name to help prevent a problem that might not occur for days or weeks. During every online/offline of the resource and every one hour that the name is online, Windows Server 2016 Technical Preview checks for the following:

- A searchable domain controller
- A synchronized Cluster Name Object (CNO) password
- A CNO in Active Directory that is turned on
- An existing CNO and Virtual Computer Object (VCO) in Active Directory

Windows Server 2016 Technical Preview also includes several additional tests in cluster validation for network names to do the following:

- Check that the CNO and VCO are greater than 15 characters
- Verify that the CNO has Create Computer Object permissions in the Organizational Unit (OU) in the Active Directory of which it is a member
- Ensure that it is possible to sign in to the CNO and corresponding VCO
- Confirm that the local Users group on the nodes has the members CLISUR and NT AUTHORITY\Authenticated Users

The previous items are commonly the cause of issues with network names, which is why these new diagnostics were added.

Cluster Operating System Rolling Upgrade

Windows Server 2016 Technical Preview introduces a wonderful new method for upgrading the operating system of your server clusters with no downtime and dramatically reduced effort. This feature is called the *Cluster Operating System Rolling Upgrade*. Cluster operating system rolling upgrades initially will be limited to Hyper-V clusters and scale-out file servers (SOFS) clusters for Windows Server 2016 Technical Preview.

Until now, the cluster administrator was tasked with developing a detailed migration plan to update clusters with a new operating system. Often, administrators waited to move a cluster until new hardware was brought in as part of a system refresh. This often meant several years without any new capabilities for the cluster and some planned downtime for moving services between the old and new cluster.

Cluster Operating System Rolling Upgrade does not require the purchase of any additional hardware; the upgrades are done in place to each of the nodes. The cluster itself never needs to be stopped or restarted; the work takes place at the cluster node level, and all services remain online during the rolling upgrade process. Unlike typical cluster migration strategies, you do not need to make a new cluster. The existing cluster objects, including cluster name and cluster IPs, remain the same and online during the upgrade. Even better news is that you can fully reverse the process until the Cluster Functional Level attribute is changed. (More on that later.)

Figure 2-12 shows a three-node Hyper-V cluster named ContosoPVTCloud.




Figure 2-12: Failover Cluster Manager console with a three-node cluster

Figure 2-13 and Figure 2-14 show three cluster shared volumes and a couple of VMs running.




Figure 2-13: Failover Cluster Manager console with two VMs running under Roles




Figure 2-14: Failover Cluster Manager console with three cluster shared volumes

As you would expect, you can move all services and roles between all nodes in the cluster. It is a fully functional and updated cluster. This is a good time to take backups, especially of the cluster itself if you choose to restore prior to completing the upgrade process.

At this point, you can begin the rolling operating system upgrades of the cluster. This is a test environment with few services running. In your environment, ensure that you have enough cluster resources to allow for one node at a time to be upgraded with the workloads supported by the remaining nodes active in the cluster.

Begin by evicting one of the nodes from the cluster to start the rolling upgrade process. You can pause and evict the node from within Failover Cluster Manager or, in Windows PowerShell, you can run the `Suspend-ClusterNode` cmdlet, followed by the `Remove-ClusterNode` cmdlet. You can choose any node in the cluster to begin the rolling upgrade. This example begins with ClusterHost01 of the ContosoPVTCloud cluster. First, right-click the node and then, on the shortcut menu, select `Pause` to pause it, as illustrated in Figure 2-15.




Figure 2-15: Pausing Cluster Node

Next, right-click the node and select Evict to evict it, as depicted in Figure 2-16.




Figure 2-16: Evicting Cluster node

Begin to install a new operating system installation on that node.

Note You do not do an in-place operating system upgrade; you format and install a new operating system by using Windows Deployment Services, install from media, or other method of your choice to deploy the new operating system.

After Windows Server 2016 Technical Preview has been installed on ClusterHost01, install the Hyper-V role, Failover Clustering, and, if necessary, Multipath I/O. Configure the network and storage as it was configured prior to the reinstallation. This is a good time to check for any updates available for this version of Windows Server 2016 Technical Preview. Join ClusterHost01 to the ContosoCloud domain that contains the cluster being upgraded.

Sign in to ClusterHost01 as a domain administrator of ContosoCloud domain or as another user with permissions on the current ContosoPVTCloud cluster. Launch Failover Cluster Manager and click Add Node, or run the Add-ClusterNode cmdlet. This must be done on the Windows Server 2016 server Technical Preview, not from a current cluster member running Windows Server 2012 R2.

The Failover Cluster Manager console on ClusterHost01 shows that it has successfully joined the cluster, as shown in Figure 2-17.

Nodes (3)				
Search				
		Queries		
Name	Status	Assigned Vote	Current Vote	Information
ClusterHost01	Up	1	1	
ClusterHost02	Up	1	1	
CLUSTERHOST03	Up	1	1	

Figure 2-17: Cluster nodes successfully joined

Both VM roles have been moved to ClusterHost01, as demonstrated in Figure 2-18.

Roles (2)					
Search					
		Queries			
Name	Status	Type	Owner Node	Priority	Information
VDI_WIN81	Running	Virtual Machine	ClusterHost01	Medium	
Win2012R2_VM01	Running	Virtual Machine	ClusterHost01	Medium	

Figure 2-18: VM roles have been moved to ClusterHost01

One of the cluster shared volumes and the quorum disk witness has also been moved to ClusterHost01, as shown in Figure 2-19.

Disks (4)					
Search					
		Queries			
Name	Status	Assigned To	Owner Node	Disk Number	Cap
Cluster Disk 1	Online	Disk Witness in Quorum	ClusterHost01	5	
Cluster Disk 2	Online	Cluster Shared Volume	ClusterHost01	4	
Cluster Disk 3	Online	Cluster Shared Volume	CLUSTERHOST03	3	
Cluster Disk 4	Online	Cluster Shared Volume	CLUSTERHOST03	2	

Figure 2-19: Drives moved to ClusterHost01

You can move all roles and resources between any nodes in the cluster. This is not a one-way move to the new cluster. All nodes function in the cluster normally and can host any role or resource in the cluster during this mixed operating system phase of the rolling update.

While the cluster is mixed between Windows Server 2012 R2 and Windows Server 2016 Technical Preview, you can patch and maintain all nodes normally until the rolling upgrade is completed. Backups can occur, but you should exclude them on the node being upgraded.

Continue the cluster operating system rolling upgrade by following the same process on ClusterHost02 and ClusterHost03. Again, drain and evict one node at a time, rebuild that node with Windows Server 2016 Technical Preview, and then rejoin to the domain and the cluster.

To perform the same steps by using Windows PowerShell, use the following examples:

1. Pause one of the nodes and drain off the roles, for example as follows:

```
PS C:\> Suspend-ClusterNode -Drain -TargetNode 2012R2-NODE4
```

2. Evict the node from the cluster by using the following command:

```
PS C:\> Remove-ClusterNode -Name 2012R2-NODE4
```

3. Perform a clean installation of Windows Server Technical Preview to the node that was evicted.
4. Add the Failover Clustering feature by using the following command:

```
PS C:\> Install-WindowsFeature -ComputerName 2012R2-NODE4 -Name Failover-Clustering -IncludeManagementTools -IncludeAllSubFeature
```

5. Add the Windows Server Technical Preview node to the Windows Server 2012 R2 cluster by using the following command:

```
PS C:\> Add-ClusterNode -Cluster Cluster -Name Preview-Node5
```

```
PS C:\> Get-ClusterNode
```

Name	ID	State
----	--	-----
2012R2-NODE3	1	Up
PREVIEW-NODE5	2	Up

6. Reinstall roles, features, and software being used on the cluster (Hyper-V, SQL, etc.).
7. Test failovers.
8. If everything proceeded properly, repeat steps 2 through 7 on the remaining node(s).

While some nodes are running Windows Server 2012 R2 and the others are running Windows Server 2016 Technical Preview, you will be running in a mixed operating system mode. Associated with this mode is your cluster functional level. While you are in this mixed operating system mode of your cluster, running the following Windows PowerShell command from a Windows Server 2016 Technical Preview node returns a value of 8 for Windows Server 2012 R2 and a value of 9 for Windows Server 2016 Technical Preview nodes:

```
Get-ClusterNode | ft -auto NodeName, MajorVersion, MinorVersion, BuildNumber, NodeHighestVersion,
@{Expression={$_.NodeHighestVersion -shr 16}; Label="NHV.Cluster Functional Level";width=21},
@{Expression={$_.NodeHighestVersion -band 0xffff};Label="NHV.Cluster Upgrade Version";width=24},
NodeLowestVersion,
@{Expression={$_.NodeLowestVersion -shr 16}; Label="NLV.Cluster Functional Level";width=21},
@{Expression={$_.NodeLowestVersion -band 0xffff};Label="NLV.Cluster Upgrade Version";width=24}
```

For example, Figure 2-20 shows the Windows PowerShell output with just one node remaining to be upgraded.

NodeName	MajorVersion	MinorVersion	BuildNumber	NodeHighestVersion	NHV.Cluster Functional Level	NHV.Cluster Upgrade Version	NodeLowestVersion
ClusterHost01	6	4	9881	589827	9	3	524291
ClusterHost02	6	4	9881	589827	9	3	524291
CLUSTERHOST03	6	3	9600	333888	8	9600	333888

Figure 2-20: Nodes that can be upgraded

The output indicates the NodeHighestVersion and the NodeLowestVersion. When these two values match on all nodes, you can upgrade the cluster functional level by running the new cmdlet, Update-ClusterFunctionalLevel, as presented in Figure 2-21.

```
PS C:\Users\administrator.CONTOSOCLLOUD>
PS C:\Users\administrator.CONTOSOCLLOUD> Update-ClusterFunctionalLevel
Updating the functional level for cluster ContosoPVTCloud.
Warning: You cannot undo this operation. Do you want to continue?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): Y
Name
----
ContosoPVTCloud
PS C:\Users\administrator.CONTOSOCLLOUD>
```

Figure 2-21: Updating the functional cluster level

When all nodes in the cluster have been upgraded to Windows Server 2016 Technical Preview, run the Windows PowerShell command again, and this time you will get a value of 9, indicating that all nodes are upgraded, as depicted in Figure 2-22.

```
PS C:\Users\administrator.CONTOSOCLLOUD> Get-ClusterNode | ft -auto NodeName, MajorVersion, MinorVersion, BuildNumber, NodeHighestVersion, NHV.Cluster Functional Level, NHV.Cluster Upgrade Version
NodeName      MajorVersion MinorVersion BuildNumber NodeHighestVersion NHV.Cluster Functional Level NHV.Cluster Upgrade Version
-----
ClusterHost01 6             4           9881        589827            9                             3
ClusterHost02 6             4           9881        589827            9                             3
ClusterHost03 6             4           9881        589827            9                             3
```

Figure 2-22: All nodes upgraded

After this process is complete, your cluster operating system has been upgraded on all nodes and is fully functional for running Hyper-V and SOFS clusters. During this process, there was no downtime; all services were available. Because there are VMs and file shares that are configured with functionality from Windows Server 2012 R2, there are still a few steps on the cluster to get all of the functionality of Windows Server 2016 Technical Preview.

If you look at the VMs in Hyper-V Manager, you can see that the version number of the VM is 5.0, as illustrated in Figure 2-23.




Figure 2-23: VM Version number

This indicates Windows Server 2012 R2 configuration for that VM. To upgrade the Hyper-V configuration version for each VM during your next maintenance window, you can run the Get-VM cmdlet followed by Update-VMConfigurationVersion, as shown in Figure 2-24.

```
PS C:\Users\administrator.CONTOSOCLLOUD> Get-VM
Name      State CPUUsage(%) MemoryAssigned(M) Uptime      Status
-----
VDI_WIN81 Off    0           0               00:00:00    Operating normally
Win2012R2_VM01 Off    0           0               00:00:00    Operating normally

PS C:\Users\administrator.CONTOSOCLLOUD> Update-VMConfigurationVersion
cmdlet Update-VMConfigurationVersion at command pipeline position 1
Supply values for the following parameters:
Name[0]: VDI_WIN81
Name[1]: Win2012R2_VM01
Name[2]:

Confirm
Are you sure you want to perform this action?
Performing a configuration version update of "VDI_WIN81" will prevent it from being migrated to or imported on previous versions of Windows. This operation is not reversible.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): y

Confirm
Are you sure you want to perform this action?
Performing a configuration version update of "Win2012R2_VM01" will prevent it from being migrated to or imported on previous versions of Windows. This operation is not reversible.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): y
PS C:\Users\administrator.CONTOSOCLLOUD>
```

Figure 2-24: Updating VM Configuration

This will change the VM configuration version to 6.1 and turn on all new features available for Hyper-V in Windows Server 2016 Technical Preview, as shown on the Summary tab of the VM depicted in Figure 2-25.




Figure 2-25: VM Configuration Version

For SOFS servers that have completed the cluster operating system rolling upgrade process, you will also want to upgrade the functionality of the storage pools. They will still be limited to the functionality in Windows Server 2012 R2 until you run the Update-StoragePool cmdlet. Specify the name of the storage pool that you want to update in this process.

When you have completed the cluster operating system rolling upgrade on all nodes, updated the cluster functional level, and updated the VM configuration version and storage pool to Windows Server 2016 Technical Preview functionality, you are all done!

Gone are the days when you needed to create a new cluster and recreate or move workloads. Cluster Operating System Rolling Update simplifies long-term cluster management and makes those complexities a thing of the past.

Hyper-V improvements

By Robert Mitchell, Deepak Srivastava, Shabbir Ahmed, and Ramnish Singh

Microsoft Hyper-V virtualization technology has been enhanced in a number of ways in Windows Server 2016 Technical Preview, and this section describes several of these improvements. Robert Mitchell demonstrates a new feature called Virtual Machine Groups and also describes the new cross-version VM mobility capabilities of the platform. Deepak Srivastava walks you through the new VM configuration version, new configuration file format, and new support for using checkpoints in production environments. Finally, Shabbir Ahmed and Ramnish Singh demonstrate the new hot add and remove capability for network adapters and memory that is now supported by the Hyper-V role.

VM groups

To make the management of multiple VMs easier, Windows Server 2016 Technical Preview has added support for groupings of VMs. VM groups are exactly what the name implies: logical groupings of VMs.

There are two different types of groups:

- VM collections
- Management collections

A *VM collection* group is a logical collection of VMs. This type of group makes it possible for administrators to carry out their tasks on specific groups, rather than having to carry them out on each individual VM separately.

A *management collection* group is a logical collection of VM collection groups. With this type of group, administrators can nest VM collections as needed.

In Hyper-V Manager, it is possible to carry out operations on multiple VMs simply by selecting multiple objects, as illustrated in Figure 2-26.




Figure 2-26: Options available on VM

You can carry out these tasks without using VM groups, but the effort is somewhat limited. You can do more by using VM groups. Two scenarios for which VM groups are useful are backups and VM replicas. Even though it is fairly easy to back up or replicate a VM, and although such functionality has been included in Windows Server for some time, all VMs are dealt with separately. In some situations, because of distributed applications, VMs should be treated as a unit. This is true in both backup and VM replica situations.

Creating VM collections

The following new Windows PowerShell cmdlets have been introduced to facilitate scripting:

- New-VMGroup
- Get-VMGroup
- Remove-VMGroup
- Add-VMGroupMember
- Remove-VMGroupMember
- Rename-VMGroup

As of this writing, VM group management tools are still being developed; however, they will be visible in Windows PowerShell, Hyper-V Manager, and the upcoming version of Microsoft System Center Virtual Machine Manager.

To group together the three example VMs shown in Figure 2-27, you need to do the following:

1. Create a VM group.
2. Add the VMs to the group membership.




Figure 2-27: VM Groups

The code that follows is a Windows PowerShell script that will accomplish our goals. Keep in mind that the VM group being created is a VM collection group. Only VM collection groups can have VMs directly placed within them.

```
#Setup VM variables
$VM1 = Get-VM -Name VM1
$VM2 = Get-VM -Name VM2
$VM3 = Get-VM -Name VM3

#Create new VM Group
New-VMGroup -Name TestVMG1 -GroupType VMCollectionType

#Setup VM Group variable
$TestVMG1 = Get-VMGroup -Name TestVMG1

#Add VMs to the group/collection
Add-VMGroupMember -VMGroup $TestVMG1 -VM $VM1
Add-VMGroupMember -VMGroup $TestVMG1 -VM $VM2
Add-VMGroupMember -VMGroup $TestVMG1 -VM $VM3
```

The result of these steps is a VM group that contains three VMs.

This can be verified by using the management tools and querying either the VMs or the VM groups. The following example shows how to do this by utilizing the Get-VM and Get-VMGroup cmdlets, respectively:

```
PS C:\> Get-VM | ft Name, state, groups - AutoSize

Name State Groups
----
VM1 Running {TestVMG1}
VM2 Running {TestVMG1}
VM3 Running {TestVMG1}

PS C:\> Get-VMGroup * | ft Name, vmmembers -AutoSize

Name VMMembers
----
TestVMG1 {VM2, VM3, VM1}
```

The updated Get-VM cmdlet lists what groups (if any) of which the VM is a member. A VM can be a member of multiple groups. If this is the case, the Get-VM cmdlet will return a list of multiple groups.

The new Get-VMGroup lists any VMs that are members of a specified group, or, as in the preceding example, in which we use a wildcard, all existing groups. In the example, we query all groups because we know there is just one. However, we can add one of the VMs to the membership of second group. Here is a quick Windows PowerShell script that will do just that:

```
#Create new VM Group
New-VMGroup -Name TestVMG2 -GroupType VMCollectionType

#Setup VM Group variable
$TestVMG2 = Get-VMGroup -Name TestVMG2

#Add VMs to the group
Add-VMGroupMember -VMGroup $TestVMG2 -VM $VM1
```

Using the Get-VM cmdlet, we can see that VM1 now belongs to both the TestVMG1 group and the new TestVMG2 group:

```
PS C:\> Get-VM | ft Name, state, groups - AutoSize

Name State Groups
----
VM1 Running {TestVMG2, TestVMG1}
VM2 Running {TestVMG1}
VM3 Running {TestVMG1}
```

Using the Get-VMGroup cmdlet, we now see both groups and VM1 are members of both VM groups:

```
PS C:\> Get-VMGroup * | ft Name, vmmembers -AutoSize
```

Name	VMMembers
TestVMG2	{VM1}
TestVMG1	{VM2, VM3, VM1}

There are now two VM groups: one comprising three VMs, and the other with a single VM, as shown in Figure 2-28.




Figure 2-28: Multiple VM groups

With the two VM groups established, you can carry out actions directed at VM1, VM2, and VM3 by utilizing TestVMG1. You can perform actions directed only at VM1 by utilizing TestVMG2.

Creating management collections

VM collections are fairly simple. They maintain a membership of VMs. Management collections, on the other hand, maintain a membership of VM collections. Figure 2-29 shows a management group that contains both of the VM groups that were created earlier. Those VM groups contain actual VMs. Note that VMs cannot directly belong to the membership of a management collection.




Figure 2-29: Single management group containing multiple VM Groups

Creating management groups is nearly identical to creating VM groups using the management tools previously outlined. The following Windows PowerShell script creates a new management group and adds both of the existing VM groups to it:

```
#Create new Management Group
New-VMGroup -Name TestVMGM1 -GroupType ManagementCollectionType

#Setup Management Group variable
$TestVMGM1 = Get-VMGroup -Name TestVMGM1

#Add VM Groups to the Management Group
Add-VMGroupMember -VMGroup $TestVMGM1 -VMGroupMember $TestVMG1
Add-VMGroupMember -VMGroup $TestVMGM1 -VMGroupMember $TestVMG2
```

An interesting difference between VM groups and management groups is that management groups can contain both VM groups and other management groups. Put simply, this means that you can nest management groups.

The following Windows PowerShell script creates a new management group named Outside and adds our first management group, TestVMGM1, to its membership:

```
#Create new Management Group
New-VMGroup -Name OutsideGroup -GroupType ManagementCollectionType

#Setup Management Group variable
$OutsideGroup = Get-VMGroup -Name OutsideGroup

#Add VM groups to the Management Group
Add-VMGroupMember -VMGroup $OutsideGroup -VMGroupMember $TestVMGM1
```

The management group (OutsideGroup) contains another management group (TestVMGM1), which contains the two VM groups (TestVMG1 and TestVMG2), which contain different groupings of three VMs (VM1, VM2, and VM3), as demonstrated in Figure 2-30.




Figure 2-30: Multitier management groups

Finally, you can use the previously described management tools to determine which VMs and which groups are members of other groups.

Obviously, this nesting capability opens an entirely new dimension in how you can organize VMs. VMs become objects that you can group much like user and computer objects in Active Directory. This will be more visible when you use this capability in conjunction with the upcoming version of Virtual Machine Manager.

True VM mobility

Being able to move VMs from one host to another has been a must since the inception of Hyper-V. In the early days of Hyper-V, during the Windows Server 2008 timeframe, only offline migration was possible (see Figure 2-31). The VM was taken offline, moved, and then brought back online. This was done by using the export and import functionality. Although this offered some VM mobility, it was restrictive in that it required downtime for the VM.




Figure 2-31: Offline migration

In Windows Server 2008 Hyper-V, you could move a VM from one host to another host only when the VM was offline.

Later, with the release of Windows Server 2008 R2, live migration first made it possible to move a VM while it was still running. However, live migration was only available between clustered Hyper-V hosts where the VMs lived on a cluster shared volume (CSV), as shown in Figure 2-32.




Figure 2-32: Live migration

Windows Server 2008 R2 Hyper-V introduced the ability to move running VMs from one cluster node to another cluster node.

A completely new level of freedom came with Windows Server 2012 and its ability to live-migrate VMs between any Hyper-V hosts of the same version (see Figure 2-33), regardless of whether either the source or destination was part of a failover cluster.




Figure 2-33: Any host same OS live migration

Windows Server 2012 Hyper-V introduced the ability to move running VMs from any host to any other host.

Windows Server 2012 R2 took live migration a step further, introducing the first “cross version” live migration. VMs could live-migrate from any Windows Server 2012 host to any Windows Server 2012 R2 host, regardless of its membership in a failover cluster (see Figure 2-34).




Figure 2-34: 2012 to 2012 R2 live migration

Windows Server 2012 R2 Hyper-V introduced the ability to move running VMs from a host running Windows Server 2012 to a host running Windows Server 2012 R2.

Windows Server 2016 Technical Preview breaks yet another boundary with down-level migration, giving administrators true freedom of control over their VMs. Previously, live migration would work only between hosts running the same version of Windows Server or the next version of Windows Server. The table that follows summarizes the migration options available for Hyper-V in each version of Windows Server running on the host:

Host operating system	Migration options
Windows Server 2008	Offline migration only
Windows Server 2008 R2	Live migration only between cluster nodes
Windows Server 2012	Live migration into or out of cluster
Windows Server 2012 R2	Live migration into or out of cluster, and from down-level Windows Server
Windows Server 2016 Technical Preview	Live migration into or out of cluster, and to up-level or down-level Windows Server

Windows Server 2016 Technical Preview is the only version that gives you the ability to live-migrate to a host running an earlier version of Windows Server (see Figure 2-35).




Figure 2-35: Migration from 2016 to earlier version of windows

Windows Server 2016 Technical Preview Hyper-V introduces the ability to move running VMs to a host running an earlier version of Windows Server.

For VMs on Windows Server 2016 Technical Preview to live-migrate to earlier versions of Windows Server, the following must be true:

- Both hosts must be members of the same Active Directory.
- Both hosts must have live-migration functionality turned on.

Turning on live migration has not changed from previous versions. On the host device, go to the Hyper-V Settings dialog box and select the Enable Incoming And Outgoing Live Migrations option, and then select from where you would like to receive incoming live migrations, as shown in Figure 2-36.




Figure 2-36: Live migration settings for a host

The mechanics of performing a live migration are the same as they were in previous versions of Windows Server. There are three ways to carry out the process:

- Use Hyper-V Manager on the host
- Create a script in Windows PowerShell
- Use Virtual Machine Manager (not included as part of Windows Server)

When using Hyper-V Manager, right-click the VM that you want to migrate and then, on the shortcut menu, select Move, as shown in Figure 2-37.




Figure 2-37: VM shortcut menu

To do the same operation using Windows PowerShell, use the Move-VM cmdlet. The following example moves a VM named New Test VM to a destination server named Hyper-Server:

```
PS C:\> Move-VM "New Test VM" Hyper-Server
```

Note The preceding cmdlet moves the VM to the Hyper-V host's default location.

Keep in mind that even though any VM can live-migrate from Windows Server 2012 to any newer Windows Server host, only version 5.0 VMs can migrate from Windows Server 2016 Technical Preview down to Windows Server 2012 R2. You can view the version in Hyper-V Manager (shown in Figure 2-38) or by utilizing the Get-VM cmdlet in Windows PowerShell.




Figure 2-38: VM version number

Note Do not confuse version with generation. Both Generation 1 and Generation 2 can be version 5.0. The version number has to do with the version of Windows Server that was used to create the VM, whereas the generation has to do with what virtualized hardware is available to the VM.

It is also important to note that although you can live-migrate VMs outside of failover clustering, you will most likely use this new mobility within failover clustering. For the first time since Windows Server 2003, failover clustering now supports mixed mode clusters. This means that you can upgrade Windows Server 2012 R2 cluster nodes to the new Windows Server 2016 Technical Preview while retaining their cluster membership. And with the improvements to mobility, you can move VMs effortlessly between older and newer cluster nodes as part of the overall cluster upgrade strategy.

VM configuration version

The VM upgrade process has changed in Windows Server 2016 Technical Preview. In the past, when you imported VMs to a new version of Hyper-V, they were automatically upgraded. However, it was not always easy to identify which VMs were imported from a previous version of Hyper-V and which were newly created. That's because the VM configuration version upgrades automatically with the host upgrade.

The real challenge, however, was that you couldn't roll back the VM to a previous version of the VM configuration. The VM configuration version determines with which versions of Hyper-V the VM's configuration, saved state, and snapshot files are compatible. In Windows Server 2016 Technical Preview, the VM configuration version upgrade process is no longer automatic. This makes it possible for you to move the VM to a server running an earlier version of Hyper-V, such as Windows Server 2012 R2. In that case, you do not have access to new VM features until you manually update the VM configuration version.

All VM capabilities remain compatible such as live migration, storage live migration, and dynamic memory. Hence, upgrading a VM is now a manual operation that is separate from upgrading the physical host. It is important to note that when you upgrade the configuration version of the VM, you cannot downgrade it. If you use VMs that were created with Windows Server 2012 R2, you will not have access to new VM features until you manually update the VM configuration version.

VMs with configuration version 5.0 are compatible with Windows Server 2012 R2 and can run on both Windows Server 2012 R2 and Windows Server 2016 Technical Preview. VMs with configuration version 6.0 are compatible with Windows Server 2016 Technical Preview but will not run on Hyper-V running on Windows Server 2012 R2.

Upgrading the configuration version

To upgrade the configuration version, shut down the VM and, at an elevated Windows PowerShell command prompt, type the following command:

```
Update-VmConfigurationVersion vmname or vmobject.
```

To check the configuration version of the VMs running on Hyper-V, from an elevated command prompt, run the following command:

```
Get-VM * | Format-Table Name, Version
```

To illustrate the configuration version upgrade process, the following example determines the VM configuration version imported from a host running Windows Server 2012 R2 and then shows how to upgrade its configuration version. In this case, as expected, the configuration version of the VM is 5.0 as indicated in Hyper-V Manager (see Figure 2-39).



Figure 2-39: VM version number

You can confirm this by using Windows PowerShell as follows:

```
PS C:\Users\Administrator> Get-VM vm02 | Format-Table Name, Version
Name                                     Version
----                                     -
vm02                                     5.0
```

As stated previously, you must shut down the VM and run the following Windows PowerShell command to upgrade the configuration version of the VM:

```
PS C:\Users\Administrator> Update-VMConfigurationVersion vm02
Confirm
Are you sure you want to perform this action?
Performing a configuration version update of "vm02" will prevent it from being migrated to or imported on
previous versions of Windows. This operation is not reversible.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): Y
PS C:\Users\Administrator>
```

When checked again, the configuration version in Hyper-V Manager now has the value 6.0, as depicted in Figure 2-40.



Figure 2-40: Upgraded version number

Again, you can confirm this by using Windows PowerShell as follows:

```
PS C:\Users\Administrator> Get-VM vm02 | Format-Table Name, Version
Name
----
vm02
Version
-----
6.0
```

As an aside, if you get any startup failure after a VM configuration version upgrade, try turning on secure boot and then run the following Windows PowerShell command:

```
Set-VMFirmware -VMName "VMName" -SecureBootTemplate MicrosoftWindows
```

The VM configuration version is successfully upgraded, which means the VM has access to new VM features introduced in Windows Server 2016 Technical Preview.

Upgrade process considerations

You need to be aware of several considerations before you upgrade the configuration version of a VM:

- You must shut down the VM before you upgrade the VM configuration version.
- The configuration version upgrade process is one way; that is, when you upgrade the configuration version of the VM from version 5.0 to version 6.0, you cannot downgrade, and, hence, afterward you cannot move the VM to a server running Windows Server 2012 R2.
- The Update-VMConfigurationVersion cmdlet is blocked on a Hyper-V cluster when the cluster functional level is Windows Server 2012 R2. You can still move the VM between all of the nodes in the Hyper-V cluster, however, when the cluster has a mix of both Windows Server 2012 R2 and Windows Server 2016 Technical Preview.

New configuration file format

After you have upgraded the VM configuration version as described in the previous section, the VM will use the new configuration file format. The new VM configuration file format uses the .vmcx extension for the VM's configuration data and the .vmrs extension for its runtime state data. The new format is a binary file format, which means that you cannot edit the file directly. The new configuration file format increases the efficiency of reading and writing the VM's configuration data, reduces the potential for data corruption in the event of a storage failure, and provides better overall efficiency.

Figure 2-41 shows the new VM configuration file format, which uses the .vmcx extension for the VM's configuration data and the .vmrs extension for runtime state data.

Name	Date modified	Type	Size
 EAF3B45D-6929-43A2-82E1-05A65F31A6CC	11/4/2014 3:48 AM	File folder	
 EAF3B45D-6929-43A2-82E1-05A65F31A6CC.vmcx	11/6/2014 2:49 AM	VMCX File	95 KB
 EAF3B45D-6929-43A2-82E1-05A65F31A6CC.VMRS	11/6/2014 2:49 AM	VMRS File	4,194,380 KB

Figure 2-41: VM configuration files

You can determine a VM's configuration location and related information by using Windows PowerShell to examine the properties of the VM:

```
PS C:\Users\Administrator> Get-VM -Name vm02 |Format-List *
VMName           : vm02
VMId             : eaf3b45d-6929-43a2-82e1-05a65f31a6cc
Id              : eaf3b45d-6929-43a2-82e1-05a65f31a6cc
Name            : vm02
State           : Running
IntegrationServicesState : Update required
OperationalStatus : {Ok}
PrimaryOperationalStatus : Ok
SecondaryOperationalStatus :
StatusDescriptions : {Operating normally}
PrimaryStatusDescription : Operating normally
SecondaryStatusDescription :
Status          : Operating normally
Heartbeat       : OkApplicationsHealthy
ReplicationState : Disabled
ReplicationHealth : NotApplicable
ReplicationMode  : None
CPUUsage        : 0
MemoryAssigned  : 4294967296
MemoryDemand    : 600834048
MemoryStatus    :
SmartPagingFileInUse : False
Uptime         : 22:37:12
IntegrationServicesVersion : 6.3.9600.16384
ResourceMeteringEnabled : False
AutomaticCriticalErrorAction : Pause
AutomaticCriticalErrorActionTimeout : 30
ConfigurationLocation : c:\vmdata\vm02\vm02
SnapshotFileLocation : c:\vmdata\vm02\vm02
CheckpointType    : Production
AutomaticStartAction : StartIfRunning
AutomaticStopAction : Save
AutomaticStartDelay : 0
SmartPagingFilePath : c:\vmdata\vm02\vm02
NumaAligned       : True
NumaNodesCount    : 1
NumaSocketCount   : 1
Key               : Microsoft.HyperV.PowerShell.VirtualMachineObjectKey
IsDeleted         : False
ComputerName      : SIGGPB04-T1
Version          : 6.0
Notes            :
Generation       : 2
Path             : c:\vmdata\vm02\vm02
CreationTime      : 11/4/2014 3:44:13 AM
IsClustered      : False
SizeOfSystemFiles : 97132
ParentSnapshotId  :
ParentSnapshotName :
MemoryStartup     : 4294967296
DynamicMemoryEnabled : False
MemoryMinimum     : 536870912
MemoryMaximum     : 1099511627776
ProcessorCount    : 1
RemoteFxAdapter   :
NetworkAdapters   : {Network Adapter}
FibreChannelHostBusAdapters : {}
ComPort1         : Microsoft.HyperV.PowerShell.VMComPort
ComPort2         : Microsoft.HyperV.PowerShell.VMComPort
FloppyDrive       :
DVDDrives        : {}
HardDrives        : {Hard Drive on SCSI controller number 0 at location 0}
VMIntegrationService : {Time Synchronization, Heartbeat, Key-Value Pair Exchange,
Shutdown...}
```


Production checkpoints

Windows Server 2016 Technical Preview introduces a new concept of taking checkpoints for production VMs; that is, production checkpoints. A *checkpoint* is a point-in-time capture of the state of a VM, which gives you the ability to revert the VM to an earlier state. Before Windows Server 2016 Technical Preview, the use of checkpoints focused on test and development scenarios but was not recommended for use in production environments.

Production checkpoints deliver the same kind of experience as in Windows Server 2012 R2, but they are now fully supported for production environments for two main reasons:

- The Volume Snapshot Service (VSS) is now used instead of saved state to create checkpoints.
- Restoring a checkpoint is just like restoring a system backup.

Note VSS is used for creating production checkpoints only on Windows VMs; Linux VMs do this by flushing their file system buffers to create a file system-consistent checkpoint.

If you want to create checkpoints by using saved-state technology, you can still use standard checkpoints for your VM. However, the default for new VMs will be to create production checkpoints with a fallback to standard checkpoints.

In certain scenarios, an administrator might need to disable checkpoints for specific VMs for operational reasons. This is now feasible in Windows Server 2016 Technical Preview, which gives you the ability to turn on or turn off production checkpoints on individual VMs. This option provides flexibility and gives Hyper-V administrators the means to manage and optimize their resources effectively.

Figure 2-42 demonstrates how you can use VM settings to turn on or turn off checkpoints for the VM and allow production checkpoints. By default, the Enable Checkpoints option is selected and is configured to allow production checkpoints and to create standard checkpoints if it is not possible to create a production checkpoint.




Figure 2-42: Configuring production checkpoints on a VM

To create a new production checkpoint for a VM, turn on checkpoints for that VM, right-click the VM in Hyper-V Manager, and then, on the shortcut menu that appears, click Checkpoint, as shown in Figure 2-43.




Figure 2-43: Menu option for creating a new production checkpoint for a VM

Note If you turn off production checkpoints for a VM, the Checkpoint option will not appear in the shortcut menu for the VM.

When a production checkpoint is created, a message appears (see Figure 2-44), confirming that the production checkpoint has been successfully taken.




Figure 2-44: Message indicating that the production checkpoint was successfully created

And, of course, you also can do all of this by using Windows PowerShell.

Hot add and hot remove for network adapters and memory

With Windows Server 2016 Technical Preview, you no longer need to plan for downtime to upgrade or downgrade memory on VMs hosted on Hyper-V. There is also no downtime when adding or removing a network card. Now, you can hot-add and hot-remove both network adapters and memory on the platform. This is a huge improvement that will ease the Hyper-V administrator's job. In a physical environment, installing additional RAM or adding a new network card is a time-consuming process that usually involves planning and downtime. With this new feature, you can accomplish everything with no downtime. Both service providers and enterprises can now scale up or scale down the memory of VMs in seconds by using either Hyper-V Manager or Windows PowerShell.

Note Hot-add memory works for generation-1 and generation-2 guests running Windows Server 2016 Technical Preview. It does not work with Windows Server 2012 R2 or earlier.

Hot add and remove memory

Figure 2-45 presents Hyper-V Manager with two VMs named VM1 and VM2 running on it. Hyper-V Manager shows that VM2 is a generation-1 VM, and the Settings dialog box for VM2 shows that this VM has been provisioned with 2 GB (2048 MB) of RAM.




Figure 2-45: Generation-1 settings for memory

Connecting to this VM using Virtual Machine Connection shows that two applications are currently running on its desktop: Date and Time, which displays a clock with the current time, and Task Manager, which displays the memory usage of VM2 and shows that 2 GB are available, as shown in Figure 2-46.




Figure 2-46: VM task manager showing memory usage

Use the settings for VM2 to change the RAM used by this VM from 2 GB to 4 GB and then click Apply while the VM is still running. Within a few seconds, Hyper-V Manager shows that VM2 is now running with 4 GB of RAM, with no reboot necessary, as illustrated in Figure 2-47.




Figure 2-47: VM memory settings changing while running

Virtual Machine Connection shows that the clock is still running in VM2, and Task Manager displays 4 GB of memory available to the VM, using the hot-add memory feature of Windows Server 2016 Technical Preview, as shown in Figure 2-48.




Figure 2-48: VM Task Manager showing VM with new memory

Hot add and remove network adapters

Adding or removing a network adapter while the VM is running without incurring downtime works only for generation-2 VMs running either Windows or Linux. Supported Windows operating systems include Windows Server 2016 Technical Preview.

In the following example, connecting to the generation-2 VM named VM1 using Virtual Machine Connection and opening the Network Connections folder shows that this VM has only a single network connection named Ethernet 2, as depicted in Figure 2-49.




Figure 2-49: Single VM NIC

To hot-add another network adapter to this VM, In the Settings dialog box, on the Add Hardware page, select Network Adapter, as shown for VM1 in Figure 2-50.




Figure 2-50: Adding a NIC to a VM

Note If the Network Adapter option is unavailable on the Add Hardware page of the Settings dialog box, it is because the VM is Generation 1, which does not support hot add and remove network adapter functionality.

Click Apply for the changes to take effect. After a few seconds, the new network adapter is installed while the VM is still running, as shown in the Network Connections folder in Virtual Machine Connection. Figure 2-51, demonstrates that the hot-add of the network adapter is successful.




Figure 2-51: VM displaying the newly added NIC

MultiPoint Services

By Shabbir Ahmed and Ramnish Singh

This section covers Windows MultiPoint Server. This is a shared resource technology from Microsoft that delivers low-cost shared computing, primarily for education environments. As we explain, Windows Server 2016 Technical Preview will include a new MultiPoint Services role feature with which you can use the functionality of Windows MultiPoint Server in your Windows Server deployment.

Windows MultiPoint Server

Before discussing MultiPoint Services in Windows Server 2016 Technical Preview, it might be useful to talk about Windows MultiPoint Server.

Windows MultiPoint Server, built on Windows Server technology, is a Windows solution that makes it possible for multiple users, each with their own independent and familiar Windows experiences, to simultaneously share one computer.

The first version of Windows MultiPoint Server, based on Windows Server 2008 R2, was released in February 2010 and was called Windows MultiPoint Server 2010. In March 2011, Microsoft released Windows MultiPoint Server 2011, based on Windows Server 2008 R2 SP1. The most recent version, Windows MultiPoint Server 2012, was released in November 2012 and contains several new features and upgrades from previous versions.

Like any other server, you can install Windows MultiPoint Server on a physical or virtual server, and you can create users inside the server by using MultiPoint Manager. Also, like other Windows-based servers, Windows MultiPoint Server gives you the full capabilities and experience of Windows as well as access to all of the latest updates.

Each user accesses Windows MultiPoint Server through a station. There are three types of stations for users, referred to as secondary stations:

- A device consisting of a monitor, a keyboard, and a mouse directly connected to the physical Windows MultiPoint Server computer through video cables
- A device consisting of a monitor, a keyboard, and a mouse directly connected to the physical Windows MultiPoint Server computer through a USB cable
- A traditional network client, such as a desktop, laptop, zero client, or thin client, using Remote Desktop Protocol (RDP) to connect to the Windows MultiPoint Server computer

In addition to any of these secondary stations, a Windows MultiPoint Server system includes an additional monitor, keyboard, and mouse connected to the server through a video cable and used to perform administration, configuration, and troubleshooting tasks. This station is called the primary station. The primary station is the first station to start using Windows MultiPoint Server. When not being used for administrative tasks, this station can be used like any other secondary station. As with other Windows-based servers, Windows MultiPoint Server gives you the full capabilities and experience of Windows as well as access to all of the latest updates.

You can review the diagrams linked in the following list to better understand how stations are connected and how they access Windows MultiPoint Server:

- A Windows MultiPoint Server system with four direct video connections: <http://i.technet.microsoft.com/dynimg/IC404229.gif>.
- A Windows MultiPoint Server system with one direct-video-connected station and two USB zero client-connected stations: <http://i.technet.microsoft.com/dynimg/IC568782.gif>. Assume the left-side station, which is connected through direct video, is the primary station.
- A Windows MultiPoint Server system with RDP-over-LAN-connected stations: <http://i.technet.microsoft.com/dynimg/IC568781.gif>.

Windows MultiPoint Server provides a user-friendly console called MultiPoint Manager and MultiPoint Dashboard for day-to-day administration and management. Windows MultiPoint Server is also a low-cost alternative to computing scenarios in which all users need to have their own computers, because with Windows MultiPoint Server multiple users can share a single machine. As of this writing, the maximum number of users who can connect to one instance of Windows MultiPoint Server is 20; most customers run multiple Windows MultiPoint Server instances to accommodate hundreds of users.

Depending on your requirements and restrictions (such as number of stations/users, type of activities performed by users, available space to deploy Windows MultiPoint Server and stations, and so on), you can design the layout of the physical Windows MultiPoint Server computer and stations in a variety of ways. You can find sample Windows MultiPoint Server system layouts and configuration on Microsoft TechNet at <https://technet.microsoft.com/en-us/library/jj916390.aspx>.

Windows MultiPoint Server versions

Windows MultiPoint Server is available in two versions:

Standard:

- Up to 10 simultaneously connected stations
- Joining a domain is not supported
- Does not support virtualization as a host or guest operating system

Premium:

- Up to 20 simultaneously connected stations
- Joining a domain is supported
- Virtualization is supported as a host or guest operating system

Scenario for better understanding

Imagine a classroom at a university. This room must accommodate 15 students at any given time and each student needs to do the following:

- View the presentation or content of what the teacher is presenting
- Create, access, and share files
- Get instructions and messages from the instructor
- Collaborate with others
- Get an enhanced learning experience

At the same time, the instructor should be able to do following:

- Monitor the learning experience from wherever she is delivering the presentation
- See what the students are doing by accessing their desktops
- Send messages to individuals and to the class as a whole
- Allow websites
- Use remote control to assist students when they need help

The premium version of Windows MultiPoint Server with 15 stations of either USB, video cable, or LAN-based RDP client stations (or a mix of these) can meet these requirements and make it possible for the students to learn efficiently and productively at their own workstations.

Note A single monitor can be used by two students with a split screen. You can use existing monitors, keyboards, and mice, and you can use MultiPoint Connector to manage Windows 7 or Windows 8-based PCs or laptops, just like any other secondary station. For more information, go to <http://blogs.technet.com/b/multipointserver/archive/2013/01/26/managing-pcs-and-tablets-with-multipoint-server-2012.aspx>.

MultiPoint Services role in Windows Server 2016 Technical Preview

With the release of Windows Server 2016 Technical Preview, Microsoft introduces a new role (as of this writing), called *MultiPoint Services*, with which you can use the functionality of Windows MultiPoint Server in your Windows Server deployment. MultiPoint Services makes it possible for multiple users, each with their own Windows experience, to share a single computer. When users sign in to the shared computer, they see their individual desktops instead of a standard, shared desktop.

From a technical perspective, when you install the MultiPoint Services role, Remote Desktop Session Host is also installed, which makes it possible for multiple users to access a single server. The two work together to provide lower total cost of ownership as well as the ability to use orchestration tools to manage Remote Desktop Session Host sessions and other devices.

The Windows MultiPoint Server standalone product has been providing great value for educational, retail, and other institutions from the time Windows MultiPoint Server 2010 was released in February 2010. With the next version of Windows Server, this capability will be broadly available for all commercial customers as a new server role in Windows Server. This will simplify single server deployments for education and retail scenarios by integrating the Windows Multipoint Services role within Windows Server. MultiPoint Services is a basic, low-cost, single-server multiuser offering with which customers can deploy a turnkey, highly durable, easy-to-manage solution to their sites (a *site* being a location where all users are in the same physical premises).

The MultiPoint Services role is ideal for scenarios in which the following are true:

- The server and the end user access devices that are in close proximity (physically connected or connected via a LAN), such as in a classroom lab.
- You don't need Remote Desktop Broker or RD Gateway for your deployment, unlike RDS deployment.
- You don't require Active Directory as a mandate for management and configuration.

Although originally developed for education, Windows MultiPoint Server has attracted retail and small-to-medium sized businesses since its launch. Based on customer feedback, the decision was made to integrate this functionality as a role within Windows Server to make it more broadly available and to do away with the limit of 20 users per Windows MultiPoint Server.

Using the MultiPoint Services role, each user will have an independent and familiar Windows computing experience, using his or her own monitor, keyboard, and mouse connected to the host computer. Basically, the MultiPoint Services role gives multiple users the ability to simultaneously share one computer, with each user enjoying an independent Windows experience. As with Windows MultiPoint Server, the following station types will be supported by a server running the MultiPoint Services role:

- Direct-video-connected stations
- USB-zero-client-connected stations (including USB-over-Ethernet zero clients)
- RDP-over-LAN-connected stations (for rich client or thin client computers)

From a deployment and administration perspective, the following differences exist when you use RDP-over-LAN-connected stations versus direct-video or direct-USB connections:

- You're not limited to physical USB and video connection distances.
- Scaling is easy because any client on your network can potentially be used as a remote station
- You cannot use the MultiPoint Manager console to look hardware issues.
- There is no split-screen functionality.
- There is no station renaming or automatic sign-in configuration through the MultiPoint Manager console.
- You can reuse existing computer hardware as stations.

Installing the MultiPoint Services role

Figure 2-52 shows a hypothetical lab environment created to demonstrate, step by step, how to install and work with the MultiPoint Services role in Windows Server 2016 Technical Preview.




Figure 2-52 Lab environment using Windows Server 2016 Technical Preview with the MultiPoint Services role installed with RDP-over-LAN-connected stations

With the upcoming version of Windows Server, you can install the MultiPoint Services role by using the Add Roles and Features Wizard in Server Manager. To do so, select role-based or feature-based installation and then select the MultiPoint Services role. Along with MultiPoint Services, this also installs other Windows Server role services and features that MultiPoint Services uses, such as Remote Desktop Session Host and Desktop Experience.

To install the MultiPoint Services role, start the Add Roles And Features Wizard on a Windows Server 2016 Technical Preview system, as depicted in Figure 2-53.




Figure 2-53: Installing Multipoint Services role

Next, from the server selections the wizard provides, select the server (destination) on which you want to install the MultiPoint Services role, as shown in Figure 2-54.




Figure 2-54: Select the server to which you want to install the role

This example environment has just one server named ThresholdVM; select this and then click Next. On the Server Roles page, the new option for the Multipoint Services role appears, as illustrated in Figure 2-55.




Figure 2-55: Selecting Multipoint Services role

As soon as you select MultiPoint Services, the wizard prompts you to install additional features and services that are required for the MultiPoint Services role (see Figure 2-56). You cannot install the MultiPoint Services role unless these additional services and features are also installed.




Figure 2-56: Required dependencies for the MultiPoint Services role

More info Before reading further, you might want to review the following documentation on Microsoft TechNet to understand more about these additional services and features that are installed:

- Remote Desktop Services: <http://technet.microsoft.com/library/dn283323.aspx>
- Desktop Experience Overview: <http://technet.microsoft.com/library/cc772567.aspx>
- Print and Document Services: <http://technet.microsoft.com/library/hh831468.aspx>
- Search: <http://msdn.microsoft.com/library/windows/desktop/aa965362%28v=vs.85%29.aspx>
- File and Storage Services: <http://technet.microsoft.com/library/hh831487.aspx>

Note Active Directory is not a prerequisite for installing MultiPoint Services.

The Install MultiPoint Services page opens next, as shown in Figure 2-57.




Figure 2-57: Description of multipoint services role

As shown in Figure 2-58, the MultiPoint Services role requires that you install Print And Document Services on the server.




Figure 2-58: The Print And Document Services page of the Add Roles And Features Wizard

Only the Print Server role service is installed for Print And Document Services, as illustrated in Figure 2-59.




Figure 2-59: Role Services for Print And Document Services

The MultiPoint Services role also requires Remote Desktop Services, as shown in Figure 2-60.




Figure 2-60: Remote Desktop Services

You then need to select the Remote Desktop Licensing and Remote Desktop Session Host role services (see Figure 2-61), both of which are required.




Figure 2-61: Selecting remote desktop required services

When the Confirmation page appears, click Next to install all of the roles and services required for the MultiPoint Services role. The server restarts to finish the installation.

MultiPoint Manager and MultiPoint Dashboard

Like Windows MultiPoint Server, the MultiPoint Services role features MultiPoint Manager and MultiPoint Dashboard for management and administration.

MultiPoint Manager

Figure 2-62 shows what MultiPoint Manager looks like in the example lab environment.




Figure 2-62: MultiPoint Manager

Following is a summary of the kinds of tasks you can perform by using MultiPoint Manager:

- Manage your Windows MultiPoint Server system and system tasks:
 - Edit server settings
 - Restart the computer
 - Shut down the computer
 - Switch to console mode
 - Turn on disk protection
 - Add client access licenses
 - Remove client access licenses
 - Remap all stations
 - Save connection settings to file
 - Add or remove Windows MultiPoint Server servers
 - Add or remove personal computers

- Manage user stations and station hardware:
 - Identify stations
 - Suspend all stations
 - Sign off of all stations
 - Start identifying all stations
 - Stop identifying all stations
- Manage user accounts:
 - Change full name
 - Change password
 - Change level of access
 - Add user account
- Manage virtual desktops
- Manage user files
- Troubleshooting

For different levels of access, the MultiPoint Services role features three types of users:

- **Standard user** Standard users can sign in from any station.
- **MultiPoint Dashboard user** MultiPoint Dashboard users can sign in from any station and use the MultiPoint Dashboard to manage standard user sessions.
- **Administrative user** Administrators have complete access to Windows MultiPoint Server and can make any changes on the computer.

More info For more details concerning virtual desktops, see the blog post at <http://blogs.technet.com/b/multipointserver/archive/2012/10/23/windows-multipoint-server-2012-creating-virtual-desktop-stations-part-1.aspx>.

MultiPoint Dashboard

Figure 2-63 illustrates what MultiPoint Dashboard looks like in the sample lab environment.




Figure 2-63: MultiPoint Dashboard

Following is a summary of the kinds of tasks you can perform by using MultiPoint Dashboard:

Manage user desktops:

- Project desktops:
 - Project your desktop to all desktops
 - Project your desktop to selected desktops
 - Project selected desktop to all desktops
 - Stop projecting desktop
- Manage desktops:
 - Block/unblock selected desktops
 - Block/unblock all desktops
 - Set message
- Limit web access:
 - Allow/disallow sites from list
 - Limit/stop limiting web access on selected desktops
 - Limit/stop limiting web access on all desktops
- Start applications:
 - Start applications on all desktops
 - Start applications on selected desktops
 - Close opened applications from desktops

- Instant message:
 - Send instant messages
 - Dismiss instant messages

Manage Windows MultiPoint Server systems by using MultiPoint Dashboard:

- Restart system
- Shut down system
- Remap system
- Turn on/off instant messaging on system
- Turn on/off auto launch

Common MultiPoint Services role usage capabilities

The MultiPoint Services role in Windows Server 2016 Technical Preview delivers the most important elements of the Windows client desktop experience to individual user desktops. Following are the most common MultiPoint Services role usage capabilities:

- A single computer is shared between multiple users.
- Applications are installed only on the central server, and all users and stations access these installed applications on the central server.
- Each user has a personal computing experience and private folders with no need for a separate computer.
- Multiple Windows MultiPoint Server systems are managed in a computer lab, classroom, training center, or small business environment.
- Desktop activity is monitored via thumbnail views of each user.
- Screens are blocked or unblocked with a customizable message to get the group's attention.
- Group can be restricted to only accessing one or more websites.
- One screen can be projected to the other screens to present and to demonstrate a particular task.
- Files can be created, accessed, and shared.
- Instruction and messages can be received from the primary station.
- Group can collaborate.
- Learning experience is enhanced.
- The secondary stations' learning experience can be monitored from the primary station.
- It's possible to access users' desktops to monitor what they are doing by accessing their desktops.
- Messages can be sent to individuals or all stations.

- Remote control can be used to assist a user connected to secondary station.
- It's possible to communicate privately with a standard user who asks for help.
- Tasks can be demonstrated by taking control of a standard user's keyboard and mouse.
- All of the preceding can be done with a traditional PC, laptop, or tablet device as a station using MultiPoint Connector.

Note Depending on the application scenario (productivity, mixes, video intensive, and so on), hardware sizing of the server running the MultiPoint Services role might need to be changed to accommodate the required number of users.

Hear about it first.



Get the latest news from Microsoft Press sent to your inbox.

- New and upcoming books
- Special offers
- Free eBooks
- How-to articles

Sign up today at MicrosoftPressStore.com/Newsletters



Storage

By Ned Pyle, Claus Joergensen, and Matt Garson

The introduction of Windows Server 2016 Technical Preview brings with it several new storage capabilities to the platform. References to Windows Server 2016 Technical Preview in this chapter refer to the TP4 release. This chapter describes four of these improvements with technical information provided by insiders at Microsoft. Ned Pyle provides some up-to-date information about Storage Replica that goes beyond that provided in *Windows Server Technical Preview Guide for Storage Replica* (available on Microsoft TechNet at <https://technet.microsoft.com/library/mt126104.aspx>). Claus Joergensen and Matt Garson examine Storage Spaces Shared Nothing and how you can use it to build highly available storage systems using only local storage instead of drives that are physically connected to all storage nodes.

Storage Replica

Storage Replica is a new feature in Windows Server 2016 Technical Preview with which you can set up storage-agnostic, block-level, synchronous replication between clusters or servers for disaster recovery. You also can use it to stretch a failover cluster across sites for high availability. Synchronous replication makes it possible for you to mirror data in physical sites with crash-consistent volumes, ensuring zero data loss at the file-system level. Asynchronous replication gives you the ability to extend sites beyond metropolitan ranges, albeit with the possibility of data loss. Storage Replica does not operate at a file level as does DFS Replication. Instead, it replicates data blocks and is therefore immune to issues of file locks, open handles, and so on.

More info Because Storage Replica is a new feature in Windows Server 2016 Technical Preview and the software-defined datacenter, there are no previous enhancements. For a full step-by-step scenario guide to configuring Storage Replica, search for "Storage Replica" on Microsoft TechNet.

Synchronous replication

Synchronous replication guarantees that the application writes data to at least two locations at the same time before completion of the write operation. This replication is most suitable for mission-critical data because it requires network and storage investments and carries a risk of degraded application performance. Synchronous replication is suitable for both high availability (HA) and disaster recovery (DR) solutions.

As Figure 3-1 illustrates, when application writes occur on the source data copy (1), the originating storage does not acknowledge the I/O immediately. Instead, those data changes replicate to the remote destination copy (2) and log data is written (3). The remote site then returns an acknowledgment (4). Only then does the application receive the I/O acknowledgment (5). This ensures constant synchronization of the remote site with the source site, in effect extending storage I/O across the network. In the event of a source site failure, applications can failover to the remote site and resume their operations with assurance of zero data loss.




Figure 3-1: Synchronous replication performed by Storage Replica

Note In the diagram, T indicates data flushed to the volume at the source site, and T1 indicates data flushed to the volume at the remote site. In all cases, logs always write through.

Asynchronous replication

In contrast to synchronous replication, asynchronous replication means that when the application writes data, that data replicates to the remote site without immediate acknowledgment guarantees (see Figure 3-2). This mode facilitates faster response time to the application as well as a DR solution that works geographically. With its higher-than-zero recovery-point objective (RPO), asynchronous replication is less suitable for HA solutions such as failover clusters because they are designed for continuous operation with redundancy and no loss of data.




Figure 3-2 Asynchronous replication performed by Storage Replica

As Figure 3-2 illustrates, when the application writes data (1), the replication engine captures and logs the write (2) and immediately acknowledges to the application (3). The captured data then replicates to the remote location (4). The remote node processes the copy of the data, writes log data (5), and lazily acknowledges back to the source copy (6). Because replication performance is no longer in the application I/O path, the remote site's responsiveness and distance are less important factors. There is risk of data loss if the source data is lost while the destination copy of the data is still in buffer without leaving the source.

Implementation-specific details

Storage Replica utilizes SMB 3.0 as a reliable, high-speed data transport for replication. This grants all the advantages of SMB 3.0, such as multichannel, remote direct memory access (RDMA), encryption, signing, and Kerberos-based security. Storage Replica does not require any changes to Active Directory Domain Services (AD DS) or any domain administrative permissions. The following table summarizes the implementation-specific details of Storage Replica:

Feature	Details
Type	Host-based
Synchronous	Yes
Asynchronous	Yes (server-to-server only)
Storage hardware agnostic	Yes
Replication unit	Volume (partition)
Windows Server Stretch Cluster creation	Yes
Server-to-server replication	Yes
Transport	SMB 3.0
Network	TCP/IP or RDMA
RDMA	iWARP, InfiniBand*
Replication network port firewall requirements	Single IANA port (TCP 445 or 5445)
Multipath/Multichannel	Yes (SMB 3.0)
Kerberos support	Yes (SMB 3.0)
Over the wire encryption and signing	Yes (SMB 3.0)
Per volume failovers allowed	Yes
Management UI in box	Windows PowerShell, Failover Cluster Manager
*Subject to further testing. InfiniBand might require additional long-haul equipment	

Note In Windows Server 2016 Technical Preview, Storage Replica does not implement transitive replication, the so-called “A-B-C” topology (that is, synchronous replication from server A to server B, and then asynchronous replication from server B to server C). Storage Replica does not implement one-to-many replication. It is possible, for virtualized workloads only, to use Hyper-V Replica as the secondary asynchronous replication mechanism. This means configuring Hyper-V Replica on the source A volume and replicating to a server other than B, forming an “A-to-B + A-to-C” topology.

Requirements

The following are prerequisites for Storage Replica testing:

- Windows Server 2016 Technical Preview Datacenter Edition
- At least 4 GB of physical memory in each server and at least two cores
- AD DS (for SMB to use Kerberos)
There is no need for schema updates, AD DS objects, certain AD DS functional levels, and so on.
- Network
 - Greater than or equal to 1 Gbps network between servers
 - Firewall ports open: SMB, WS-MAN
- Storage
 - One NTFS/ReFS-formatted volume dedicated to replication per server/cluster site with at least 8 GB of free space
 - GUID partition table (GPT) (not master boot record [MBR])
 - JBOD, iSCSI, local direct-attached storage (DAS) (non-cluster) SCSI or SATA, Storage Spaces Direct (cluster-to-cluster only), or Storage Array Network (SAN)
 - Same sector sizes for the data volume drives and for the log volume drives
 - No %SystemRoot% or page file located on replicated volumes or log volumes

Recommendations

The following are highly recommended for Storage Replica testing:

- Network
 - Bandwidth: ≥ 10 Gbps network between servers
 - Latency: ≤ 5 ms round-trip average for synchronous replication
- Storage
 - Flash (solid-state drive [SSD]) disks for the log volume(s) with at least 8 GB of free space

Note The final version of Storage Replica will contain additional scenarios and features, such as management by Microsoft Azure Site Recovery, and write-ordered consistency groups to ensure protection of applications that use multiple replicated volumes.

Note You can use the Test-SRTopology Windows PowerShell cmdlet to ensure that you meet the requirements and assist with recommended configuration tuning for log files.

Scenarios

Storage Replica was designed with two scenarios in mind:

- Stretching of a failover cluster for high availability
- Replication between servers for disaster recovery

Stretch cluster replication

A *stretch cluster* (also referred to as a *multisite cluster*) uses Storage Replica to connect two sets of asymmetric shared storage within a single failover cluster. This storage can be serially attached SCSI JBOD (just a bunch of drives), iSCSI target, or SAN. Cluster nodes attach to each of the two sets of storage, ostensibly in two physical locations, such as different buildings on the same campus or different metropolitan datacenters. The replicated storage can be either cluster shared volumes (CSV) or role-assigned physical disk resources (PDR).

Figure 3-3 presents the typical architecture used for implementing stretch cluster replication using Storage Replica. On the left is the Redmond site, where there are two servers (SR-SRV-01 and SR-SRV-02) and shared storage (SAN, JBOD, or iSCSI). On the right is the Bellevue site, where there are two more servers (SR-SRV-03 and SR-SRV-04) and more shared storage. You can use Storage Replica to combine the servers and shared storage at these two sites into a single stretched cluster by asymmetrically replicating storage from one site to the other.




Figure 3-3: Typical architecture used for implementing stretch cluster replication

This configuration makes a failover cluster tolerant not just of node failures, but entire site failures. When a single node in a site fails, another node in that site becomes the new source of replication. When all nodes in a site fail, a node in the other site becomes the source of replication. All of this occurs automatically, just like a normal nonstretched cluster. Stretch clustering requires a minimum of two nodes, and the cluster can contain up to 64 nodes.

In Windows Server 2016 Technical Preview, the two cluster roles recommended for replication are Hyper-V and General Use File Server. You should avoid configuring Scale-Out File Server as a stretch cluster because Windows Server failover clusters are not inherently site aware and applications will end up connecting to nodes in both sites and then redirecting back to the owning node where I/O writes occur. This potentially can lead to poor application performance. Microsoft supports the use of virtual machine (VM) guest clusters in the Technical Preview for evaluation purposes only.

You can manage this cluster with Failover Cluster Manager (cluadmin.msc) through a simple wizard-driven interface. To create a stretched cluster, simply create a CSV and configure the General Use File Server role or a Hyper-V VM role. Right-click the source storage (shown in the Figure 3-4 as Cluster Disk 3 in the cluster named np-sr-cluster.com), click Replication, and then, in the shortcut menu that appears, click Enable.




Figure 3-4: Turning on drive replication in Cluster Manager

In the Configure Storage Replica Wizard that opens, on the Select Destination Data Disk page, select a destination drive from the available storage displayed for the source drive that you want to replicate. In Figure 3-5, Cluster Disk 1 is selected as the destination data drive.




Figure 3-5: Selecting a destination drive for the replica

Follow the remaining wizard prompts to finish configuring your stretched cluster. When configured, the storage synchronously replicates between the source and destination storage on the cluster. When completed, replication forms the stretch cluster and Storage Replica protects the data on the source and destination drives. For example, Figure 3-6 shows the direction of replication after a subsequent failover. In this case, Cluster Disk 1 is now the source of replication and Cluster Disk 3 is the destination.




Figure 3-6: Disk information after a replica has been created

Note You can also use Windows PowerShell, using the Failover-Clustering and StorageReplica modules to create a stretch cluster.

Server-to-server and cluster-to-cluster replication

Storage Replica can connect two individual servers—sometimes called stand-alone replication—and their volumes. This storage can be serially attached SCSI JBOD, iSCSI target, SAN, or even local DAS, such as SCSI drives attached to a local RAID controller. Storage Replica can also replicate between two clusters as if they were two servers, with any shared storage the cluster considers acceptable. Replication occurs between two physical locations, such as different buildings on the same campus or different metropolitan datacenters. The replicated storage must be either NTFS or ReFS volumes.

Figure 3-7 depicts a typical architecture used for implementing server-to-server replication using Storage Replica. On the left is the Redmond site with server SR-SRV01 and some storage (SAN, JBOD, or iSCSI). On the right is the Bellevue site with server SR-SRV02 and more storage. You can use Storage Replica to combine the servers and storage at these two sites into a partnership by asymmetrically replicating storage from one site to the other.




Figure 3-7: Typical architecture used for implementing server-to-server replication

In the server-to-server and cluster-to-cluster scenario, there is no graphical interface and no automatic failover management—all administration is manual and human-driven through the Windows PowerShell StorageReplica module. To ensure ease of provisioning, Storage Replica implements a simple system for configuring replication with a single command when possible.

The Windows PowerShell StorageReplica module contains the following commands in Windows Server 2016 Technical Preview:

```
Get-Command -Module StorageReplica | FT -Auto
```

CommandType	Name	Version	Source
Function	Get-SRGroup	1.0	StorageReplica
Function	Get-SRPartnership	1.0	StorageReplica
Function	New-SRGroup	1.0	StorageReplica
Function	New-SRPartnership	1.0	StorageReplica
Function	Remove-SRGroup	1.0	StorageReplica
Function	Remove-SRPartnership	1.0	StorageReplica
Function	Set-SRGroup	1.0	StorageReplica
Function	Set-SRPartnership	1.0	StorageReplica
Function	Suspend-SRGroup	1.0	StorageReplica
Function	Sync-SRGroup	1.0	StorageReplica
Function	Test-SRTopology	1.0	StorageReplica

Configuring replication is as simple as providing the following information:

```
New-SRPartnership -SourceComputerName np-sr-srv05 -SourceRGName rg01
-SourceVolumeName g: -SourceLogVolumeName h: -DestinationComputerName np-sr-srv06
-DestinationRGName rg02 -DestinationVolumeName g: -DestinationLogVolumeName h:
-LogSizeInBytes 16GB
```

There are many options for the New-SRPartnership cmdlet, including creating asynchronous replication. You can also create replication in a more granular fashion by running New-SRGroup on each server and tying them together by using New-SRPartnership. You can add additional volumes to a replication group by using Set-SRGroup, and you can run more than one replication group on a server at a time. Storage Replica will include more cmdlets before the final release, including a cmdlet to determine how well replication will perform between two servers over a given network, how to optimally size the replication logs, and what the current write I/O load is on a server you propose for replication—all without the need to install or configure Storage Replica beforehand.

Storage Replica in Windows Server 2016 Technical Preview

The following are some of the key things to know concerning Storage Replica as of the Windows Server 2016 Technical Preview release:

- **Network bandwidth and latency with fastest storage** There are physical limitations to synchronous replication. Because Storage Replica implements an I/O filtering mechanism using logs and requiring network roundtrips, synchronous replication is likely to make application writes

slower. By using low-latency, high-bandwidth networks as well as high-throughput drive subsystems for the logs, you can minimize performance overhead.

- **The destination volume is not accessible while replicating** When you configure replication, the destination volume will dismount and no longer be visible in any normal GUI tools or accessible to any writes by users until you remove replication or the volume becomes the source due to failover.

Block-level replication technologies are incompatible with allowing access to the destination's mounted file system in a volume; NTFS and ReFS do not support users writing data to the volume while blocks change underneath them.

- **The Microsoft implementation of asynchronous replication is different than most** Most industry implementations of asynchronous replication rely on snapshot-based replication, whereby periodic differential transfers move to the other node and merge. In contrast, Storage Replica asynchronous replication operates just like synchronous replication, except that it removes the requirement for a serialized synchronous acknowledgment from the destination. This means that Storage Replica theoretically has a lower RPO as it continuously replicates. However, this also means it relies on internal application consistency guarantees rather than using snapshots to force consistency in application files. Storage Replica guarantees crash consistency in all replication modes.

- **Storage Replica is not DFSR** Volume-level block storage replication is not a good candidate for use in branch-office scenarios. Branch-office networks tend to be highly latent, highly utilized, and lower bandwidth, which makes synchronous replication impractical. A branch office often replicates data in a one-to-many with read-only destinations, such as for software distribution, and Storage Replica is not capable of this in the first release. When replicating data from a branch office to a main office, Storage Replica dismounts the destination volume to prevent direct access.

It is important to note, nevertheless, that many customers use Distributed File System Replication (DFSR) as a DR solution even though it is often impractical for that scenario—DFSR cannot replicate open files and is designed to minimize bandwidth usage at the expense of performance, leading to large recovery-point deltas. Storage Replica might make it possible for you to retire DFSR from some of these types of DR duties.

- **Storage Replica is not backup** Some IT environments deploy replication systems as backup solutions due to their zero-data-loss options when compared to daily backups. Storage Replica replicates all changes to all blocks of data on the volume, regardless of the change type. If a user deletes all data from a volume, Storage Replica replicates the deletion instantly to the other volume, irrevocably removing the data from both servers.
- **Storage Replica is not Hyper-V Replica or SQL AlwaysOn** Storage Replica is a general purpose, storage-agnostic engine. By definition, it cannot tailor its behavior as ideally as application-level replication. This might lead to specific feature gaps that encourage you to deploy or remain on specific application replication technologies.
- **Test only** You cannot deploy Storage Replica in production environments by using Windows Server 2016 Technical Preview. This version is only for evaluation purposes in a test lab environment.
- **Performance** The Windows Server 2016 Technical Preview version of Storage Replica is not fully performance optimized.

Storage Spaces Direct

Enterprises and service providers are continually looking for ways to deploy private clouds while reducing storage cost, accelerating initial deployment, and simplifying management and monitoring. Storage Spaces, a software-defined storage (SDS) capability introduced previously in Windows Server 2012 and enhanced in Windows Server 2012 R2, provides cost-effective shared storage solutions that are highly available, scalable, and simple to operate, providing a platform for private clouds or private hosted cloud solutions.

Windows Server 2016 Technical Preview builds on the SDS capabilities of the previous versions by introducing Storage Spaces Direct, which makes it possible for you to build highly available storage systems using local storage. This provides increased flexibility because you can use new device types such as SATA or NVME drives for creating pooled storage. Scalability is increased because you can add new server nodes with internal storage on an as-needed basis; you can use more drive devices in a single pool as well, simplifying management. Storage expandability is simplified through the action of automatic storage rebalancing. With all of these enhancements, Storage Spaces Direct provides businesses with greater flexibility for deploying storage in private cloud solutions.

To help understand Storage Spaces Direct, let's begin by examining Storage Spaces in Windows Server 2012 R2 HA storage systems. In Windows Server 2012 R2, an HA system using Storage Spaces requires drive devices to be physically connected to all storage nodes. For the drive devices to be physically connected to all storage nodes, they need to reside in an external JBOD chassis, with each storage node having physical connectivity to the external JBOD. Also, because multiple storage nodes will be connecting to each drive, the drive devices need to be serial-attached SCSI (SAS) because the SAS protocol allows for this sharing where drives such as SATA do not allow multi-initiator. Because of these requirements, this deployment is referred to as *Storage Spaces Shared JBOD* to contrast it with Storage Spaces Direct. Figure 3-8 shows a Storage Spaces Shared JBOD deployment.




Figure 3-8: Example of a Storage Spaces Shared JBOD deployment

Storage Spaces Shared JBOD provides many benefits compared to past HA storage systems. However, requiring that drive devices are physically connected to every single node limits the type of drive devices that you can use and can lead to complex SAS fabric configurations, especially as these deployments scale out.

With Windows Server 2016 Technical Preview Storage Spaces, you can now build HA storage systems using storage nodes with only local storage, which is either drive devices that are internal to each storage node or drive devices in JBODs, where each JBOD is connected only to a single storage node. This completely eliminates the SAS fabric and its complexities, but makes possible the use of drive devices such as SATA drive devices, which can further reduce cost or improve performance. Figure 3-9 presents a Storage Spaces Direct deployment.




Figure 3-9: Example of a Storage Spaces Direct deployment

It is important to also understand that Storage Spaces Direct is an evolution of Storage Spaces, which means that it is an extension of the existing SDS stack for Windows Server. Another important aspect is that Storage Spaces Direct uses SMB 3.0 for all intranode (also called *east-west*) communication, and takes advantage of all the powerful features of SMB 3.0, such as SMB Direct (RDMA-enabled NICs) for high bandwidth and low latency communication, and SMB Multichannel for bandwidth aggregation and network fault tolerance.

Implementation details

Storage Spaces Direct seamlessly integrates with the features you know today that make up the Windows Server SDS stack, including Scale-Out File Server (SMB 3.0), Clustered Shared Volume File System (CSVFS), Storage Spaces, and Failover Clustering. Figure 3-10 illustrates the Storage Spaces Direct stack.




Figure 3-10: Storage Spaces Direct stack

The updated stack includes the following, starting from the bottom:

- **Hardware** The storage system consists of a minimum of four storage nodes with local storage. Each storage node can have internal drives or drives in an external SAS-connected JBOD enclosure. The drive devices can be SATA disks or SAS disks.
- **Software Storage Bus** The Software Storage Bus spans all of the storage nodes and brings together the local storage in each node, so all drives are visible to the Storage Spaces layer above.

- **Storage Spaces** Storage Spaces provides storage pools and virtual drives. The storage pool can span all of the local storage across the nodes. The virtual drives provide resiliency to drive or node failures because data copies are stored on different storage nodes.
- **ReFS** Resilient File System (ReFS) provides the file system in which the Hyper-V VM files are stored. ReFS is a premier file system in Windows Server 2016 Technical Preview for virtualized deployments and includes optimizations for Storage Spaces such as error detection and automatic correction. In addition, ReFS provides accelerations for VHD(X) operations such as fixed VHD(X) creation, dynamic VHD(X) growth, and VHD(X) merge. CSVFS layers above ReFS bring all the mounted volumes into a single namespace.
- **Scale-Out File Server** This is the top layer of the storage stack that provides remote access to the storage system by using the SMB 3.0 access protocol.

Improved scalability

You can deploy Storage Spaces Direct using storage nodes with either local storage or nonshared JBODs. In previous versions of Windows Server, scaling out Storage Spaces solutions required a concurrent increase in the scale of the SAS fabric that joined the storage nodes to the shared SAS JBODs. In contrast, with Storage Spaces Direct, you can set up a model that removes the complexities of the SAS fabric, making scale-out as simple as adding a new storage node, either with internal storage or attached to a nonshared JBOD. Scaling out by adding storage nodes provides more flexibility in storage planning because storage expansion is no longer bound by the number of drive slots in a shared SAS JBOD.

To support this model of just-in-time scale-out, Storage Spaces Direct improves scalability compared to previous versions of Windows Server because you can now manage more drive devices in a single storage pool. Increasing the number of drive devices in a single pool reduces the number of storage pools that you must create, simplifying management of the storage solution.

Storage Spaces optimized pool

Storage Spaces Direct can optimize a storage pool to balance data equally across the set of physical drives that comprise the pool. Over time, as physical drives are added or removed or as data is written or deleted, the distribution of data among the set of physical drives that comprise the pool can become uneven. In some cases, this might result in certain physical drives becoming full, whereas other drives in the same pool have much lower consumption.

Similarly, if you add new storage to the pool, optimizing the existing data to utilize the new storage results in better storage efficiency across the pool and, potentially, improved performance from the newly available additional physical storage throughput. Optimizing the pool is a maintenance task that the administrator performs.

When the optimize pool command is started, Storage Spaces Direct moves data among the physical drives in the pool. The data movement is a background operation, designed to minimize impact to foreground or tenant workloads.

Failure scenarios

Storage Spaces Direct addresses various failure scenarios. To understand how this works, you first need to review some basic information about virtual drives.

A virtual drive consists of extents, each of which is 1 GB in size. A 100-GB virtual drive will therefore consist of 100 1-GB extents. If the virtual drive is mirrored (using `ResiliencySettingName`), there are multiple copies of the extent. The number of copies of the extent (obtained by using `NumberOfDataCopies`) can be two or three. For example, a mirrored virtual drive with three data

copies consumes 300 extents. The placement of extents is governed by the fault domain, which in Storage Spaces Direct is nodes (StorageScaleUnit), so, as shown in Figure 3-11, the three copies of an extent (A) are placed on three different storage nodes; for example, nodes 1, 2, and 3 in the figure. Another extent (B) of the same virtual drive might have its three copies placed on different nodes; for example, nodes 1, 3, and 4, and so on. This means that a virtual drive has its extents distributed across all storage nodes and the copies of each extent are placed on different nodes. Figure 3-11 depicts a four-node deployment with a mirrored virtual drive with three copies and an example layout of extents.




Figure 3-11: A four-node deployment

Next, let's take a look at various failure scenarios and examine how Storage Spaces handles them.

Scenario 1: One or more sectors on a drive has failed

In this scenario, Storage Spaces will reallocate the extent that is affected by the failing sectors. The destination drive for the reallocation could be another drive in the same node or another drive in another node that does not already have a copy of the extent. So, if the three copies of the extent are on node A, B, and C, and the extent on node A is affected by a sector failure, the new copy can be generated on a different drive in node A or any drive in Node D. Drives in node B and C cannot be used as these two nodes already have a copy of the extent.

Scenario 2: A drive has failed

In this scenario, Storage Spaces retires the physical drive from the storage pool when it discovers the drive has failed. After the physical drive has been retired, each virtual drive starts its repair process. Because the physical drive has been retired, the virtual drives generate a new copy of the extents that were on the retired physical drive. The new copies follow the same logic as in scenario 1.

Scenario 3: A drive is missing

In this scenario, Storage Spaces will do one of two things:

- If only the physical drive is missing, Storage Spaces will retire the disk.
- If the storage node or the physical enclosure to which the physical drive is attached is also missing, Storage Spaces will not retire the physical drive.

The reason Storage Spaces won't retire the physical drive in this second case is that during a storage node restart or temporary maintenance of a storage node, all the drives and physical enclosures associated with that node will be reported missing. Automatically retiring all of those drives and enclosures would potentially result in a massive amount of repair activity because you would need to rebuild all of extents on those drives elsewhere in the storage system. This could easily be multiple terabytes of data. If the drives and enclosures are really missing and will not come back to the storage system, the administrator will need to retire the missing physical drives and start the repair process.

Scenario 4: Storage node restart or maintenance

In this scenario, Storage Spaces does not automatically retire physical drives from the storage pool for the reasons described earlier in scenario 3. When the storage node comes back online, Storage Spaces automatically updates all extents that are not up to date with the copies that were unaffected by the restart or maintenance.

Scenario 5: Permanent storage node failure

In this scenario, Storage Spaces requires the administrator to retire all of the affected physical drives from the storage pool, add additional storage nodes to the storage system if needed, and then begin repair. This is not an automatic process because Storage Spaces does not know if the failure is temporary or permanent. It is not desired to initiate a repair that could potentially result in significant repair activity.

Learn more

You can find more details about Storage Spaces Direct in Windows Server 2016 Technical Preview from Microsoft TechNet (<https://technet.microsoft.com/library/mt126109.aspx>). The guide includes a detailed walkthrough using Windows PowerShell showing you how to install and configure Storage Spaces Direct in the Windows Server 2016 Technical Preview.

Deduplication

In Windows Server 2016 Technical Preview Deduplication, the core focus is to have significant impact on the scale and performance that you can address with this technology. With this shift in focus, you can now use Windows Server 2016 Technical Preview in the following scenarios:

- Volumes up to 64 TB
- File sizes up to 1 TB
- Virtualized backup
- Nano server support
- Rolling cluster upgrades

Since deduplication was introduced in Windows Server 2012, the core principles for getting the data to a chunked state have remained the same, now let us address what has changed that makes these new scenarios possible.

The optimization engine has been upgraded from single thread to multiparallel thread using multiple I/O streams, as shown in Figure 3-12.




Figure 3-12: Single-thread processing versus parallel processing

On top of this parallel execution, the algorithm for processing files has been redesigned using a new stream map structure and improved partial file optimization. This accommodates the scalability and performance to handle files up to the 1 TB size limit. Figure 3-13 shows how the mapping technology has changed to allow for better optimization overall on the volume.




Figure 3-13: Old mapping versus new stream map structures in Windows Server 2016

Although Microsoft technically supported and provided guidance to use deduplicated volumes with DPM in Windows Server 2012 R2, now further improvements to ensure ease of configuration and support are included within Windows Server 2016 Technical Preview. You can configure this directly in the Windows 2016 GUI or via Windows PowerShell by using the Enable-DeDup cmdlet, but using the new type "Backup."

```
Enable-DedupVolume -Volume <volume> -UsageType Backup
```

Another new great option is support for rolling cluster upgrade. You can begin the process of upgrading your cluster nodes to Windows Server 2016 Technical Preview and maintain the deduplication process in Windows Server 2012. Previously this was unsupported. During the migration to Windows Server 2016, all jobs will run in Windows Server 2012 mode.

Finally, a minor update, but one still worth noting, is the availability of a Storage Management API (SMAPI) interface for use with System Center 2016 Virtual Machine Manager. This makes it possible for you to set up storage deduplication and status reporting from within System Center 2016 Virtual Machine Manager.

Storage Quality of Service

One of the "missing" features from previous editions of Windows Server was the ability to apply Quality of Service (QoS) policies in relation to storage traffic. This became a huge problem in virtualization estates; if you wanted to prioritize certain workloads and give them, for example, a guaranteed level of I/O operations per second (IOPS), you simply couldn't do it.

In Windows Server 2016 Technical Preview we can now enforce resource fairness or prioritization depending on the policies you want to configure for your storage. The core usage of storage QoS will be focused around Hyper-V VMs deployed on either a Scale-Out File Server or Hyper-V Cluster with Cluster Shared Volumes.

Figure 3-14 demonstrates that we can create various policies and apply them to our different VMs, giving you the ability to have different service standards or prioritize business-critical workloads.




Figure 3-14: Storage QoS policies and their application to different VM tiers

Storage QoS in Windows Server 2016 Technical Preview is turned on by default. This means that you don't need to install an additional role or feature to get going.

For example, Figure 3-15 illustrates that if you have a Hyper-V failover cluster, you can see a new cluster resource is listed.




Figure 3-15: Storage QoS Cluster Core Resource

You also can use the the PowerShell cmdlet `Get-ClusterResource` to display the same result.

```
Get-ClusterResource -Name "Storage Qos Resource"
```

Storage QoS will only work effectively if you decide to configure appropriate policies. You can use policies to control the traffic flow as necessary based on your requirements. You can configure Storage QoS policies on the Scale-Out File Server. You essentially have a choice of two policy types:

- **Single-instance** Using single-instance policies, you can create a minimum and maximum amount of IOPS per policy. This is aggregated against a VM. For example, if a VM has a single VHD/VHDX, it will have full use of all the IOPS in the assigned policy. However, if the VM has three VHD/VHDX and they are all assigned the same single-instance policy, that VM will share the maximum number of IOPS across all drives, degrading the overall performance. You have the option to have multiple single-instance policies and configure each drive to use a different single-instance policy to ensure that they get access to all the IOPS. If you have two VMs with a single VHD each and all assigned to the same single-instance policy, they will also share the minimum and maximum IOPS.
- **Multi-instance** With multi-instance policies, again you have options to create a minimum and maximum number of IOPS. However, in this scenario, if you had two VMs with a single VHD/VHDX each, they will get their own allocation of IOPS, both minimum and maximum. However, the same rules apply that if the VM had multiple drives: unless assigned individual policies, they will share the total amount of assigned minimum and maximum IOPS.

To create a policy, use the following Windows PowerShell cmdlet:

```
$GoldVmPolicy = New-StorageQosPolicy -Name Gold -PolicyType MultiInstance -MinimumIops 100 -MaximumIops 500
```

This sample will store information about the policy in the variable. There is one property called the `PolicyId`, which we will require. To access the `PolicyId` use the following syntax:

```
$GoldVmPolicy.PolicyId
Guid
----
Cd5f6b87-fa15-402b-3545-32c2f456f6e1
```

The Guid is what we will require in order to apply this policy to a VHD by using the following Windows PowerShell sample:

```
Get-VM -Name GoldSrv* | Get-VMHardDiskDrive | Set-VMHardDiskDrive -QoSPolicyID Cd5f6b87-fa15-402b-3545-32c2f456f6e1
```

After the policy is applied, you will, of course, want to verify that the policy is indeed active, but also you will want to monitor if it is having the appropriate effect. You can do this by by using the `Get-StorageQoSFlow` cmdlet.

The following output shows what is applied and the amount of storage the IOPS the VM is actually using:

```
Get-StorageQoSFlow -InitiatorName GoldVm1 | Format-List
```

```

FilePath:c:\ClusterStorage\Volume1\VMS\Gold\GoldVM1.V          HDX
FlowId: ebfeeb54-e47a-5a2d-8ec0-0940994ff21c
InitiatorId           : ae4e3dd0-3bde-42ef-b035-9064309e6fec
InitiatorIOPS         : 464
InitiatorLatency      : 26.2684
InitiatorName         : GoldVM1
InitiatorNodeName     : node1.contoso.com
Interval             : 300000
Limit                 : 500
PolicyId              : cd5f6b87-fa15-402b-3545-32c2f456f6e1
Reservation           : 500
Status                : Ok
StorageNodeIOPS       : 475
StorageNodeLatency    : 6.5625
StorageNodeName       : node1.contoso.com
TimeStamp             : 2/12/2016 3:28:49 AM
VolumeId              : 2d34fc5a-2b3f-9922-23f4-43563b2a6787
PSComputerName        :
MaximumIops           : 100
MinimumIops           : 500

```

You can use the Get-StorageQoSFlow cmdlet to validate before you create policies what the VMs are actually using in relation to Storage IOPS.

```
Get-StorageQoSFlow | Sort-Object StorageNodeIOPS -Descending | ft InitiatorName,
@{Expression={$_.InitiatorNodeName.Substring(0,$_.InitiatorNodeName.IndexOf('.')});Label="InitiatorNodeName"},
StorageNodeIOPS, Status, @{Expression={$_.FilePath.Substring($_.FilePath.LastIndexOf('\')+1)};Label="File"}
-AutoSize
```

InitiatorName	InitiatorNodeName	StorageNodeIOPS	Status	File
GoldVM5	node1	2482	Ok	IOMETER.VHDX
GoldVM2	node2	344	Ok	BUILDVM2.VHDX
GoldVM1	node2	597	Ok	BUILDVM1.VHDX
GoldVM4	node1	116	Ok	BUILDVM4.VHDX
GoldVM3	node2	526	Ok	BUILDVM3.VHDX
GoldVM4	node1	102	Ok	

More info You can find additional scenarios to get started with Storage QoS on Windows Server 2016 at <https://technet.microsoft.com/library/mt126108.aspx>.

Networking

By Yuri Diogenes and David Branscome

In this chapter, we look at the new networking features within Windows Server 2016 Technical Preview. Specifically, we will discuss the improvements around software-defined networking and how consistency between the Azure network concepts and implementations is being brought on premise with Windows Server 2016 Technical Preview. We will also take a look at the Web Application proxy and what has been updated in Windows Server 2016.

Software-defined network

As the evolution of software-defined datacenters (SDDCs) continues, traditional networking methods such as virtual local area networks (VLANs) begin to become difficult to manage and maintain. The requirements to centrally manage and control the network landscape directly from software to dynamically create what is needed when it is needed is a key piece to this evolving concept. Introduced in Windows Server 2012 R2, software-defined networking (SDN) has evolved further to become more Azure consistent.

In this section, we want to dive into the following areas:

- Network virtualization
- Network Controller
- Remote Access Service (RAS) Gateway Multitenant Border Gateway Protocol (BGP) Router
- Software load balancer
- Datacenter firewall

Network virtualization

For IT pros managing the IT infrastructure, efficiency in hardware resources is a major consideration in the decision-making process. In the past decade, this mindset has contributed to bringing virtualization into the mainstream enterprise and achieving the efficiency of have a one-host-to-

many-different workload approach. With networking, there are some inherent limits within the stack, such as a maximum of 4,096 VLANs. Even though this is a large number, there are segments within the industry (such as service providers) for which that number can be reached quickly. In these cases, virtualized networks can provide a solution. However, this is not the only case in which network virtualization is interesting. Take, for example, a company that expands via acquisition. In that scenario, the company might purchase companies in which IP spaces overlap the enterprise network plan or layout. What if the acquired companies have licensing agreements that are tied to the IP addresses on the servers and the agreements are not easily broken or the original company no longer exists?

Whatever the case, network virtualization provides the foundation for achieving an SDN solution for the datacenter, as demonstrated in Figure 4-1.




Figure 4-1 Network virtualization

Figure 4-1 illustrates some key concepts. The first is the *routing domain ID* or RDID. You can also consider this as the virtual network, such as in Azure. Each virtual network is the boundary of isolation in that network, meaning that only subnets within that virtual network can communicate with one another. Because a virtual network is Layer 3 (i.e., IP), it makes it easy to isolate the traffic between the different networks. The isolation is enforced through Network Virtualization Generic Route Encapsulation ID (NVGRE ID) or Virtual Extensible Local Area Network (VXLAN) Network Identifier (VNI) field. In each routing domain, we can have multiple virtual subnets (VSID), and they have the ability to communicate with one another. In Figure 4-1, RDID 1 has three subnets, each with different VSIDs. If VSID 5001 wants to communicate with VSID 5002, it is possible via Layer 2 forward within the virtual switch. When the packet traverses the switch, it becomes encapsulated and mappings are applied (encapsulation header), and then it is sent to the destination Hyper-V port or destination switch. To summarize, if traffic has the same RDID, it can be forwarded between VSIDs; if the RDIDs are different, then we need to make use of a gateway.

For network administrators who are used to the concepts such as Address Resolution Protocol (ARP) broadcast, Media Access Control (MAC) addresses, and broadcast domain, these concepts all technically still apply but now fall into specific buckets. For example, if we take VSID 5001, it can be considered to be a broadcast domain or VLAN. Now, if two machines in VSID 5001 want to communicate with each other, they look up a MAC address via an ARP query to the Hyper-V switch. The Hyper-V switch has a Hyper-V port for every network adapter in a virtual machine (VM), and we know that every network adapter has a MAC address. The Hyper-V switch will keep a lookup table or ARP table of these entries so that if an ARP query comes in, it knows to where to switch the traffic. If no entry exists, it will generate an ARP broadcast in an attempt to find what port hosts the computer

to which it wants to send the traffic. This subnet or VSID will also have an IP subnet allocated to it, something like 192.168.0.0/24, which could be a subset of the virtual network allocation of 192.168.0.0/16.

Hyper-V Network Virtualization on a single host is a relatively easy concept. However, when we introduce multiple hosts that span the datacenter and we want to honor the RDID and VSID of isolated networks and subnets across the datacenter, the topic becomes a little more complex and we must introduce some address space concepts. Figure 4-2 shows a basic concept in which we have multiple subnets that overlap each other.




Figure 4-2: Address space concepts

If we did not have these subnets isolated, we would get a lot of trouble on our networks and things certainly would not work properly. However, because network virtualization isolates the IP address space, they can coexist easily on the same physical network. First, take the Blue network in Figure 4-2 or the Red network; these address spaces signify the *Customer Address (CA) Space*. For them to communicate between hosts across a physical network, they must be encapsulated via a common IP subnet. This common subnet that connects the host is called the *Provider Address (PA) Space*. Essentially, what then happens is the CA Space is mapped to an IP address in the PA space, and when VMs allocated to the CA space want to communicate across hosts, they do so via their mapped IP in the PA space.

How this mapping happens depends on the type of technology we use. In network virtualization, we can use VXLAN or NVGRE.

The following is an extract from Technet (<https://technet.microsoft.com/en-us/library/mt238303.aspx>) detailing VXLAN and NVGRE:

The Virtual eXtensible Local Area Network (VXLAN) (RFC 7348) protocol has been widely adopted in the market place, with support from vendors like Cisco, Brocade, Arista, Dell, HP, and others. The VXLAN protocol uses UDP as the transport. The IANA-assigned UDP destination port for VXLAN is 4789 and the UDP source port should be a hash of information from the inner packet to be used for ECMP spreading. After the UDP header, a VXLAN header is appended to the packet which includes a reserved 4-byte field followed by a 3-byte field for the VXLAN Network Identifier (VNI) – VSID – followed by another reserved 1-byte field. After the VXLAN header, the original CA L2 frame (without the CA Ethernet frame FCS) is appended.

NVGRE is used as part of the tunnel header. In NVGRE, the VM's packet is encapsulated within 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 4-3.




Figure 4-3: NVGRE

Network Controller

In Windows Server 2016 Technical Preview, it is very important to note that this is a new implementation of the network virtualization approach. Let's call it V2. Why is this important, well a lot of things have fundamentally changed and in V2, being consistent to how network is structured and implemented in Azure was a principal goal. Although the common concept remains the same say for network virtualization as we generally described earlier, it means some overhauled technology in Windows Server 2016 Technical Preview. One of the first things worth mentioning before diving into Network Controller and the Software Load Balancer is that the Hyper-V extensible switch has also changed; this means that extensibility options which were implemented in V1 of network virtualization will not work on V2. V2 implements the Azure Virtual Filtering Platform technology to ensure the consistency model across the private and public cloud.

In this section, we are going to describe Network Controller, which essentially is now the brain of the virtualized network solution we will implement. In large, complex networks with traditional networking technology, we would have implemented a central tool to manage these networks. Using these traditional tools, we would provide a central point of management so that we can automate configuration, maintenance, backup, and troubleshooting of the physical switch environment. In a virtualized network environment, this is exactly what Network Controller will do.

Network Controller can interact with the network and be interacted with through two different APIs specifically for each function. The Northbound API (implemented as a REST API) is used to interact with Network Controller and monitor the network and implement configuration changes. The Southbound API is used to interact with network devices and detect service configurations and basically understand the network. Using tools such as System Center Operations Manager and System Center Virtual Machine Manager, we can manage and monitor our network directly from these consoles.

Because Network Controller is considered the brain, it can manage all the networking virtualization technologies included with Windows Server 2016 Technical Preview. The following table gives a breakdown of the areas Network Controller can manage and a description of the types of management it can do.

Component	Manageable areas
Fabric network management	- Physical fabric management - IP subnets - VLANs - Layer 2 switches - Layer 3 switches - Network adapters in hosts
Firewall management	- Manage firewall rules into the vSwitch port of the VM estate - Log traffic centrally on the switch
Network monitoring	- Monitoring physical network - Monitoring virtual network - Active networking monitoring - Element data collected using SNMP polling and traps - Impact analysis
Network topology and discovery management	Discovery of network topology through elements
Software load balancer	Manage and configure load balancing rules
Virtual network management	- Virtual network policies - All elements of network virtualization
RAS gateway management	- Add and remove gateway VMs from the cluster and specify the level of backup required - Site-to-site virtual private network (VPN) gateway connectivity between remote tenant networks and your datacenter using IPsec - Site-to-site VPN gateway connectivity between remote tenant networks and your datacenter using Generic Routing Encapsulation (GRE)

	• Point-to-site VPN gateway connectivity so that your tenants' administrators can access their resources on your datacenter from anywhere • Layer 3 forwarding capability • Border Gateway Protocol (BGP) routing, which allows you to manage the routing of network traffic between your tenants' VM networks and their remote sites
--	---

More info In this book, we cannot provide an exhaustive review of all features and details in relation to Network Controller. We would encourage you to review the public TechNet article at <https://technet.microsoft.com/library/dn859239.aspx> for additional information.

Additionally, Network Controller is a complex element to deploy and get working. The following articles will provide you with the most up to date documentation for deploying and configuring Network Controller in Windows Server 2016 Technical Preview:

"Installation and preparation requirements for deploying Network Controller":
<https://technet.microsoft.com/library/mt691521.aspx>

"Deploy Network Controller using Windows PowerShell":
<https://technet.microsoft.com/library/mt282165.aspx>

RAS Gateway Multitenant BGP router

When you deploy network virtualization and employ the encapsulation and isolation methods described earlier in this chapter, you face an interesting problem: How do the VMs in these isolated networks communicate outside the isolated network? How do external machines communicate with these isolated VMs if they needed to?

Windows Server 2016 Technical Preview introduces additional capability to the RAS Gateway role to include BGP support. We previously supported the following features for RAS Gateway:

- Site-to-site VPN
- Point-to-site VPN
- GRE tunneling
- NAT

Given that all these features are now available in RAS Gateway, you can reap the following benefits:

- VMs can talk to other networks outside the routing domain to which they are assigned.
- You can create endpoints into the virtual network if required.
- You can connect virtual and physical networks together.

With the introduction of BGP, new possibilities open up for our network environments. For example, Express route works on BGP, not to mention the Internet!

BGP dynamically learns which networks are attached and announces these networks to other BGP-capable routers. The other BGP-capable routers can populate their routing tables with the entries, and if the BGP router receives a request to send traffic to a tenant's network, it will know how to route the traffic appropriately. An important part of BGP is the ability to provide route redundancy and automatically recalculate the best route to the desired network. In that case, if you have several

routers connected together and there were multiple paths to a desired network, BGP would work out the optimal route to it. Then, in the event of a failure, it would recalculate the route and announce it back to its BGP peers.

One of the challenges in relation to virtualized networks and the RAS Gateway in Windows Server 2012 is the relationship of gateways. If you want to deploy a high-availability (HA) pool of gateways in Windows Server 2012, there is no way to separate the gateway pool for separate functions or tenants. There are also strict placement requirements of the Gateway Nodes, which cause a lot of network problems in enterprise clusters.

Windows Server 2016 Technical Preview implements a true pool model in which you can create pools for specific functions or mix the functions. Figure 4-4 shows how you can deploy pools for different functions.




Figure 4-4: RAS Gateway pools

More Info RAS Gateway Multitenant BGP Routing is a complex area for discussion and is constantly changing. To view the latest information, go to <https://technet.microsoft.com/library/mt679502.aspx>.

Software Load Balancing

Windows Server 2016 Technical Preview introduces Software Load Balancing, which provides high availability and scalability to tenant workloads. Software Load Balancing includes the following features:

- Layer 4 (L4) load balancing services for "North-South" and "East-West" TCP/UDP traffic.
- Public and internal network traffic load balancing.
- Supports dynamic IP addresses on VLANs and on virtual networks that you create by using Hyper-V Network Virtualization.
- Health probe support.

- Ready for cloud scale, including scale-out capability and scale-up capability for multiplexers and host agents.

Software Load Balancing has been specifically designed to handle throughput on a scale of tens of gigabytes per cluster, making this a viable alternative to traditional hardware load balancers.

At this point, before we dive into the Software Load Balancing, let's define a few terms:

- **Virtual IP Address** This is the IP that to which external connections will route
- **Dynamic IP Address** This is the set of IPs on the VMs backing the service

When you have a service that requires Software Load Balancing, Network Controller is notified of the request and provisions a Software Load Balancing multiplexer. You can have several different multiplexers in an environment. Each multiplexer will be assigned a virtual IP address. The BGP then announces the virtual IP address to the network. The multiplexer is also responsible for accepting connections and routing them to the VMs backing the service. Because the virtual IP address is announced through BGP and is controlled by Network Controller, in the event of a multiplexer failure, Network Controller has the ability to recover by initiating a new multiplexer and reannouncing the routes through BGP. Figure 4-5 shows the Software Load Balancing architecture.




Figure 4-5: An overview of Software Load Balancing

More info Software Load Balancing requires that Network Controller be installed and configured. For instructions on how to do this, go to view the TechNet article at <https://technet.microsoft.com/library/mt632286.aspx>.

Datcenter firewall

Introduced in Windows Server 2016 Technical Preview, the datacenter firewall is designed to be a network-layer firewall with the following features:

- Stateful packet inspection
- Multitenant
- Five-tuple rule matching (Protocol, source and destination port numbers, source and destination IP addresses)

This is a multitenant option; you can use it to protect tenant VM workloads and configure it via the tenant administrators. This means that it can implement the security policies by which your organization is governed. Figure 4-6 illustrates the datacenter firewall.




Figure 4-6: Datacenter firewall

The datacenter firewall is controlled by Network Controller. The tenant administrator can configure policies and apply them directly to a Vport on the Hyper-V switch. Additionally, as tenant workloads move around the datacenter, the policy for the tenant can follow them on their journey between hosts.

Web Application Proxy

By Yuri Diogenes and David Branscome

In this section, we demonstrate how you can use the updated Web Application Proxy in Windows Server 2016 Technical Preview to easily access information from anywhere.

Publishing capability enhancements

Users can access company data by using different form factors (e.g., laptop computers, tablets, and smartphones), which here, for simplicity, we will just refer to as *devices*. These devices can originate requests from different locations, but the users expect to have an experience similar to what they have when they are on-premises. IT must ensure that the entire communication channel is secure, from data at rest in the datacenter (on-premises or in the cloud), to data in transit until it reaches the destination device. There, it will also be at rest and must also be secure.

To make it possible for users to securely access company data, Web Application Proxy in Windows Server 2016 Technical Preview was enhanced to cover more bring-your-own-device (BYOD) scenarios, such as Pre-Auth with Microsoft Exchange Server, which we will discuss later in more detail. Web Application Proxy continues to make use of Active Directory Federation Services (AD FS) and Active Directory Domain Services (AD DS) for authentication and authorization. This integration is very important for BYOD scenarios because it provides the capability to create custom rules for users who are accessing resources while physically located on-premises versus those accessing resources via the Internet.

Note If you are not familiar with Web Application Proxy in Windows Server 2012 R2, read the article at <http://technet.microsoft.com/library/dn584107.aspx>.

The Web Application Proxy installation experience is similar to that in the previous version of Windows Server 2012 R2; therefore, you can use the same steps to install it in Windows Server 2016 Technical Preview. When the installation is complete, you are prompted to perform the post-deployment configurations, as shown in Figure 4-7.




Figure 4-7: Post-deployment configuration for Web Application Proxy

Note Before you deploy Web Application Proxy, ensure that you plan the infrastructure according to the recommendations from the article at <http://technet.microsoft.com/library/dn383648.aspx>. This article was written for Windows Server 2012 R2, but the recommendations still apply to Windows Server 2016 Technical Preview.

When you finish the post-deployment steps, which are basically connecting your Web Application Proxy server to the AD FS server, you will be able to use the Publish New Application Wizard. You will notice some changes that were introduced in this new version. The first change you will notice when you click Publish under the Web Application Proxy management tool are the options available in the left pane. For example, on the Preauthentication page, you can choose either Active Directory Federation Services (AD FS) or Pass-Through for the preauthentication method, as shown in Figure 4-8.




Figure 4-8: Preauthentication selection

For the purpose of this example, select Active Directory Federation Services (AD FS) as the preauthentication method, and then click Next. On the Supported Clients page, your options are Web And MSOFBA, HTTP Basic, or OAuth2, as depicted in Figure 4-9.




Figure 4-9: The Supported Clients page

You can use the Web And MSOFBA option to preauthenticate clients by using the Microsoft Office Forms Based Authentication (MSOFBA) protocol. MSOFBA provides form-based authentication instead of basic or NTLM authentication when you use Office client applications. The second option is the well-known HTTP basic authentication that you can use in scenarios such as Exchange Active Sync (ActiveSync). This is a new capability included in this release of Web Application Proxy. For the ActiveSync scenario, the authentication process includes four core steps:

1. Web Application Proxy stops the request and passes all credentials to AD FS.
2. AD FS validates, applies policy, and replies with a token.
3. Upon success, Web Application Proxy allows the request to pass to the Exchange server.
4. Web Application Proxy caches the token for future use.

The third option is OAuth2, which is an authorization framework that many vendors use, including Microsoft. Web Application Proxy has supported OAuth2 since Windows Server 2012 R2; however, the option was not available in the user interface (UI).

More info To learn more about OAuth2, go to <http://tools.ietf.org/html/rfc6749>. You can find additional information about AD FS support for OAuth2 at <http://technet.microsoft.com/library/dn383640.aspx>.

After you select the appropriate client for the publication, click Next. The Publishing Settings page includes one new option with which you can turn on HTTP to HTTPS redirection, as illustrated in Figure 4-10.




Figure 4-10: The Publishing Settings page

This is a great addition because to turn on HTTP to HTTPS redirection in Windows Server 2012 R2 you were required to install and configure Internet Information Services (IIS). Notice that this option is selected by default but ensure that it is selected for your app before clicking Next and moving on to the Confirmation page.

Remote Desktop Gateway scenario

The changes that were first introduced in the Windows Server 2012 R2 August 2014, update package regarding the way Web Application Proxy handles Remote Desktop Gateway (RDG) publication is incorporated in this release. This change simplifies the deployment experience for IT administrators who are planning to publish RDP via Web Application Proxy and makes it possible for RDG to pick up the session cookie that was used by Remote Desktop Web Access to authenticate the RDP over HTTP traffic.

Auditing access to resources

Windows Server 2016 Technical Preview introduces a new capability that gives IT administrators better audit access to published resources. Web Application Proxy now adds to every request an X-Forwarded-For (XFF) header to verify whether the header already exists. If so, Web Application Proxy concatenates the client IP to this header.

Note XFF is a nonstandard HTTP header that became de facto standard. It is used extensively by proxy servers to identify the IP of an originated request. For more information about this, read the RFC at <http://tools.ietf.org/html/rfc7239>.

Another important aspect of Web Application Proxy auditing capabilities are the events that are logged in the Event Viewer. In this release, the Event Viewer includes many more events, such as analytics and debug logs. You will review some examples of these events in the section "Web Application Proxy troubleshooting" later in this chapter.

Taking application proxies to the modern IT world

A few years ago, our team had a big dilemma. We had two products in the market: Forefront Threat Management Gateway and Forefront Unified Access Gateway. Both of these products had been around for many years and had been deployed by tens of thousands of customers. Both of them had evolved since they were first introduced during the 1990s.

However, both products had similar issues: They were very complex products that were hard to deploy, troubleshoot, and maintain. This was partly because over the years they accumulated many capabilities that became irrelevant. At the same time, they lacked or had limited support for modern technologies such as federation and OAuth2. On top of it all, they were expensive products that had their own licenses.

It was a tough decision, but we decided to start from a blank page, to examine all the functionality of reverse proxy, to pick and choose only the technologies that matter today, and to implement them by using a fresh code base built on the most modern standards. A big part of this decision was that we wanted to embed the reverse proxy into Windows Server. We wanted to make it just like any other role service available to install from Server Manager for us, this means adhering to the strictest standards regarding code and management. Microsoft customers expect that all Windows Server role services are managed the same way, including in Windows PowerShell, the administrator UI, the remote administrator UI, performance counters, the System Center Operations Manager pack, event logs, and so on.

This is how Web Application Proxy was born in Windows Server 2012 R2. We made no compromises on code security, management, and standardization. And, we were happy that customers got it. Companies were able to deploy and integrate Web Application Proxy into their infrastructure very easily.

The downside of this approach is that we were not able to include all of the functionality we wanted to have—functionality that would make it possible for all customers to move from Threat Management Gateway and Unified Access Gateway to the new solution. However, now that we have built a solid foundation, it is easier to add more functionality to make Web Application Proxy the obvious choice to publish on-premises resources such as Microsoft SharePoint, Lync, and Exchange to remote users. This version marks an important milestone in the journey we started quite a few years ago.

Now, it is time for us to start another journey to bring remote access to the cloud era. We have created Azure Active Directory Application Proxy as another tool for customers to publish applications in cloud-based solutions. Fortunately, Web Application Proxy in Windows Server and Azure Active Directory Application Proxy share a lot of code. More than that, they share the same concepts and perception of remote access and how to make it simple to deploy and easy to maintain.

Going forward, we will continue to develop both products. We plan to offer Microsoft customers a choice with regard to which architecture to use. The cloud offers users a unique and highly efficient way to implement remote access utilizing the rich functionality and robust security mechanisms of Azure Active Directory, without the need to change their perimeter network. The same service that takes care of 18 billion authentication requests per week handles your on-premises applications.

Meir Mendelovich, Senior Program Manager

Publishing Exchange Server 2013

As noted earlier, the retirement of Forefront Threat Management Gateway left a number of Exchange administrators in a quandary about how to publish their Exchange server to the Internet. Although large organizations can generally take advantage of an existing hardware load balancer infrastructure to accomplish this task, small- and medium-sized businesses might not have the funds or expertise to manage a load balancer. This is where the Web Application Server role can be very useful.

The basic principles for publishing Exchange 2013 Outlook Web App and the Exchange Admin Center through Web Application Proxy are outlined in detail at [http://technet.microsoft.com/library/dn635116\(v=exchg.150\).aspx](http://technet.microsoft.com/library/dn635116(v=exchg.150).aspx). However, to get a better understanding of some of the capabilities of Web Application Proxy on Windows Server 2016 Technical Preview, consider the very simple scenario illustrated in Figure 4-11.




Figure 4-11: Scenario demonstrating Web Application Proxy on Windows Server 2016 Technical Preview

In this scenario, a user on a nondomain-joined machine named CLT01 will be connecting to Outlook Web App by using the URL <https://mail.contoso.com/owa>. The user's request is sent over the external network to a Routing and Remote Access Server (RRAS01), which provides external DNS resolution for the zone contoso.com and routes traffic to the Contoso internal network for outside users. RRAS01 routes the request for <https://mail.contoso.com/owa> into Contoso's internal network to the Web Application Proxy server (WAP01), which is running Windows Server 2016 Technical Preview.

On WAP01, Outlook Web App has been published using AD FS preauthentication. In fact, a number of different Exchange services have been published using AD FS preauthentication or pass-through authentication, as demonstrated in Figure 4-12.

PUBLISHED WEB APPLICATIONS
All published web applications | 6 total

Name	External URL	Backend Server URL	Preauthentication
Autodiscover	https://mail.contoso.com/autodis...	https://mail.contoso.com/autodis...	Pass-through
ECP	https://mail.contoso.com/ecp/	https://mail.contoso.com/ecp/	AD FS
Exchange Web Services	https://mail.contoso.com/ews/	https://mail.contoso.com/ews/	Pass-through
OAB	https://mail.contoso.com/oab/	https://mail.contoso.com/oab/	Pass-through
Outlook Anywhere	https://mail.contoso.com/rpc/	https://mail.contoso.com/rpc/	Pass-through
Outlook Web App	https://mail.contoso.com/owa/	https://mail.contoso.com/owa/	AD FS

Figure 4-12: Published Web Applications

In this case, we are assuming that this is a split DNS configuration and that internal and external DNS resolve the name mail.contoso.com to different IP addresses, depending on the user's location. Thus, the External URL value and Backend Server URL value are the same, but they could be different, as we'll show later. So, when a user who is on the internal Contoso network goes to <https://mail.contoso.com/owa>, authentication takes place by using Windows Integrated Authentication. This requires that the URLs for mail.contoso.com and adfs.contoso.com are defined in the Trusted Zone for Local Intranet in Internet Explorer. If that is done, the user should be able to connect to his mailbox and not be prompted for authentication at all.

On the other hand, a user connecting from outside the corporate network is presented with a form-based authentication webpage, such as that shown in Figure 4-13, and is required to provide sign-in credentials.

Figure 4-13: Form-based Authentication Page

However, one very important service used extensively in nearly every organization is missing—Microsoft Server ActiveSync. You could define a relying party trust for ActiveSync and set it up for pass-through authentication, and this what you would have done in Windows Server 2012 R2 Web

Application Proxy. But, as noted earlier in this chapter, Web Application Proxy in Windows Server 2016 Technical Preview now supports the use of HTTP Basic clients for services such as ActiveSync that don't support redirection and that use HTTP Basic to authenticate users.

HTTP Basic is the authorization method used by many protocols, including ActiveSync, to connect rich clients, including smartphones, with an Exchange mailbox. (For more information on HTTP Basic, see RFC 2617 at <http://www.ietf.org/rfc/rfc2617.txt>.) Web Application Proxy traditionally interacts with AD FS using redirections, which is not supported on ActiveSync clients. Publishing an app by using HTTP Basic provides support for ActiveSync clients in Web Application Proxy by making it possible for the HTTP app to receive a nonclaims relying party trust for the application to AD FS.

The authentication flow for clients that use HTTP Basic is described in the following steps:

1. The user attempts to access a published web application through a telephone client.
2. The app sends an HTTPS request to the URL published by WAP.
3. If the request does not contain credentials, Web Application Proxy returns an HTTP 401 response to the app containing the URL of the authenticating AD FS server.
4. The user sends the HTTPS request to the app again with authorization set to Basic and user name and Base 64 encrypted password of the user in the www-authenticate request header.
5. Because the device cannot be redirected to AD FS, Web Application Proxy sends an authentication request to AD FS with the credentials that it has: user name, password, and, if available, device certificate. The token is acquired on behalf of the device.
6. To minimize the number of requests sent to the AD FS, Web Application Proxy validates subsequent client requests by using cached tokens for as long as the tokens are valid. Web Application Proxy periodically cleans the cache. You can view the size of the cache by using the performance counter.
7. If the token is valid, Web Application Proxy forwards the request to the server backing the service and the user is granted access to the published web application.

To do this, go back to the ADFS01 server and create a nonclaims-aware relying party for ActiveSync, as depicted in Figure 4-14.

Relying Party Trusts			
Display Name	Enabled	Type	Identifier
ActiveSync	Yes	Non-Claims-Aware	https://mail.contoso.com/Microsoft-Server-ActiveSync/
Autodiscover	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/autodiscover/
Device Registration Service	Yes	WS-Trust / SAML / WS-Federation	um.ms-dfs.adfs.contoso.com
Exchange Control Panel	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/ecp/
Exchange Web Services	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/ews/
Offline Address Book	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/oab/
Outlook Anywhere	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/rpc/
Outlook Web App	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/owa/

Figure 4-14: Relying Party Trusts

Next, go to the WAP01 server, publish the ActiveSync application by using HTTP Basic, and then select the ActiveSync nonclaims-aware relying party, as shown in Figure 4-15.



Figure 4-15: Supported Clients

Note that this relying party is visible only when a nonclaims-aware relying party has been defined on the AD FS server, as demonstrated in Figure 4-16.




Figure 4-16: Relying Party Options

Complete the remainder of the wizard, configuring the external and server URL backing the service. The end result is shown in Figure 4-17.

PUBLISHED WEB APPLICATIONS			
All published web applications 8 total			
Filter			
Name	External URL	Backend Server URL	Preauthentication
APP01	https://www.contoso.com/app01/	https://apps.contoso.com/app01/	AD FS
Autodiscover	https://mail.contoso.com/autodis...	https://mail.contoso.com/autodis...	Pass-through
ECP	https://mail.contoso.com/ecp/	https://mail.contoso.com/ecp/	AD FS
Exchange Web Services	https://mail.contoso.com/ews/	https://mail.contoso.com/ews/	Pass-through
Microsoft ActiveSync	https://mail.contoso.com/Microso...	https://mail.contoso.com/Microso...	AD FS for Rich Clients
OAB	https://mail.contoso.com/oab/	https://mail.contoso.com/oab/	Pass-through
Outlook Anywhere	https://mail.contoso.com/rpc/	https://mail.contoso.com/rpc/	Pass-through
Outlook Web App	https://mail.contoso.com/owa/	https://mail.contoso.com/owa/	AD FS

Figure 4-17: All Published Web Applications

Notice that the Microsoft ActiveSync published application uses the AD FS For Rich Clients preauthentication method.

Defining the claims

Although defining claims isn't a function of the Web Application Proxy role in Windows Server 2016 Technical Preview, it's important to understand the role that claims play in a transaction. Claims are defined in the Outlook Web App section of the Actions pane on the AD FS server, as shown in Figure 4-18.




Figure 4-18: Edit Claim Rules

Select the relying party trust that you want to define claims for, and then, in the Actions pane, click Edit Claims.

In a claims-based identity model, AD FS issues a token that contains a set of claims. Claims rules govern the decisions with regard to the claims that AD FS issues. Claim rules and all server configuration data are stored in the AD FS configuration database.

To publish Outlook Web App and the Exchange Admin Center in this example, you need to make three custom claim rules:

- Active Directory user SID
- Active Directory group SID
- Active Directory UPN

When you configure the custom claims rules, you need to use the claim rule language syntax for this rule. Specifically, for the ActiveDirectoryUserSID claim rule, use the following:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
  Issuer == "AD AUTHORITY"]
=> issue(store = "Active Directory", types = ("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"), query = ";objectSID;{0}", param = c.Value);
```

When you are finished, the resulting rule will include the claim rule name and custom rule text, as depicted in Figure 4-19.




Figure 4-19: Editing Rule

Next, configure the following ActiveDirectoryGroupSID claim rule:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",  
    Issuer == "AD AUTHORITY"]  
=> issue(store = "Active Directory", types = ("http://schemas.microsoft.com/ws/2008/06/identity/  
    claims/groupsid"), query = ";tokenGroups(SID);{0}", param = c.Value);
```

And finally, configure the following ActiveDirectoryUPN claim rule:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",  
    Issuer == "AD AUTHORITY"]  
=> issue(store = "Active Directory", types = ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"),  
    query = ";userPrincipalName;{0}", param = c.Value);
```

When you're finished, click Apply, and then OK. The transform rules display the rule names on the Issuance Transform Rules tab of the Edit Claim Rules dialog box, as shown in Figure 4-20.




Figure 4-20: Edit Claim Rules

Note You will need to do this for each of your relying party trusts.

URL hostname translation

Another thing to keep in mind is that Web Application Proxy can translate host names in URLs. For example, you might have an application that your external users access by using the URL <http://www.contoso.com/app01>, whereas your internal users get to the same app by going to <http://apps.contoso.com/app01>. This is perfectly acceptable, and Web Application Proxy can handle the difference in the URL, as illustrated by Figure 4-21, in which the external URL is <http://www.contoso.com/app01> and the Backend Server URL is <http://apps.contoso.com/app01>.

PUBLISHED WEB APPLICATIONS			
All published web applications 7 total			
Filter 			
Name	External URL	Backend Server URL	Preauthentication
APP01	https://www.contoso.com/app01/	https://apps.contoso.com/app01/	AD FS

Figure 4-21: Published Web Applications

Note You cannot change the path to be <http://www.contoso.com/app1> externally and <http://apps.contoso.com/applicationXYZ> internally.

This is valuable to know when publishing Exchange, because you can have different DNS namespaces for internal and external access. Therefore, you might want to publish a URL for Outlook Web App that is different for your internal users than the one you publish for external access. Figure 4-22 shows that Web Application Proxy can accommodate this need as long as you keep the same path name.

Figure 4-22: Publish Settings for App

To allow this type of translation, you must first get the application ID for the application for which you want to allow translation. To do this, you can use the following Windows PowerShell command:

```
Get-WebApplicationProxyApplication | Format-Table ID, Name, ExternalURL
```

Figure 4-23 presents the output that results.

ID	Name	ExternalURL
E8700EC6-9BFF-5BE8-452F-50E953E901D	Outlook Web App	https://mail.contoso.com/owa/

Figure 4-23: Output for Get-WebApplicationProxyApplication

Next, take the application ID from the output shown in Figure 4-23 and enter the following Windows PowerShell command:

```
Set-WebApplicationProxyApplication -ID <application_ID>
-DisableTranslateUrlInRequestHeaders:$false
```

Enabling AD FS for your Exchange organization

When you are configuring AD FS for claims-based authentication with Outlook Web App and the Exchange Admin Center in Exchange 2013, you must turn on AD FS for your Exchange organization. This is accomplished by using the Set-OrganizationConfig cmdlet for your organization. For the example environment in this chapter, you would need to do the following:

- Set the AD FS issuer to https://adfs.contoso.com/adfs/ls.
- Set the AD FS URIs to https://mail.contoso.com/owa and https://mail.contoso.com/ecp.
- Find the AD FS token signing certificate thumbprint by using the Windows PowerShell Get-ADFSCertificate -CertificateType "Token-signing" cmdlet on the AD FS server. Then, assign the token-signing certificate thumbprint that you found.

Using the Exchange Management Shell, type the following code:

```
Get-ADFSCertificate -CertificateType "Token-signing"
```

This will provide you with the token-signing certificate's thumbprint, on which you run the following Set-OrganizationConfig cmdlets:

```
$uris = @" https://mail.contoso.com/owa", "https://mail.contoso.com/ecp"
Set-OrganizationConfig -AdfsIssuer "https://adfs.contoso.com/adfs/ls/" -AdfsAudienceUri $uris
-AdfsSignCertificateThumbprint "1a2b3c4d5e6f7g8h9i0j11k12l13m14n15o16p17q"
```

Web Application Proxy troubleshooting

The sections that follow provide a few tips on how you can troubleshoot issues that might arise in environments where Web Application Proxy has been deployed.

Collecting information about your environment

Managing and troubleshooting Web Application Proxy servers requires a good knowledge of Windows PowerShell and the cmdlets exposed for Web Application Proxy. When you are troubleshooting a Web Application Proxy problem, first take note of any error messages that appear in the console. If there aren't any obvious errors, review the event logs. You can sign in to each server and check the event logs, but you can use Windows PowerShell to simplify the process.

For example, the following Windows PowerShell command will gather all the events that the Web Application Proxy server generated in the previous 24 hours, along with their ID, Level, and Message:

```
$yesterday = (Get-Date) - (New-TimeSpan -Day 1) ;
Get-WinEvent -FilterHashTable @{LogName='Microsoft-Windows-WebApplicationProxy/Admin'; StartTime=$yesterday}
| group -Property ID,LevelDisplayName,Message -NoElement |
sort Count, Name -Descending | ft - Name -HideTableHeaders }
```

Suppose that you see Event ID 12000 repeatedly on this specific server; however, you have a number of Web Application Proxy servers, and you want to see if they are all experiencing the same error. Run the following command to collect all of the event ID 12000s generated within the previous 10 hours for a set of Web Application Proxy servers:

```
Foreach ($Server in (gwpc).ConnectedServersName){Get-WinEvent -FilterHashTable @{LogName='Microsoft-Windows-WebApplicationProxy/Admin'; ID=12000; StartTime=(Get-Date) - (New-TimeSpan -hour 10)} -ComputerName $Server -ErrorAction SilentlyContinue | group MachineName -NoElement | ft Name -HideTableHeaders }
```

Now you have the list of all the servers experiencing the issue. For this example, let's assume that there is only one server experiencing this error.

The TechNet table of error codes can be very useful for resolving the issue (<http://technet.microsoft.com/library/dn770156.aspx>). The table on TechNet suggests checking the connectivity with AD FS for this particular Web Application Proxy server. To do so, go to `https://<FQDN_of_AD_FS_Proxy>/FederationMetadata/2007-06/FederationMetadata.xml` and ensure that there is a trust relationship between the AD FS server and the Web Application Proxy server. If this doesn't work, run the Install-WebApplicationProxy cmdlet to correct the issue.

Using the Microsoft Exchange Best Practices Analyzer

You can also run the Exchange Best Practices Analyzer on the Web Application Proxy server. You can do this via the Server Manager GUI. In the far left pane, select Local Server and then, in the middle pane, scroll down to Best Practices Analyzer, as shown in Figure 4-24.




Figure 4-24: Best Practices Analyzer in Server Manager

On the right side of the Server Manager GUI, click Tasks and then select Start BPA Scan, as depicted in Figure 4-25.




Figure 4-25: Starting a BPA Scan

You can also run the Best Practices Analyzer by using the following Windows PowerShell cmdlet:

```
Invoke-BpaModel Microsoft/Windows/RemoteAccessServer
Get-BpaResult Microsoft/Windows/RemoteAccessServer
```

In this case, there is an issue related to certificate problems (see Figure 4-26)—specifically, an error message indicating “Web Application Proxy could not publish an application due to certificate problems.”




Figure 4-26: Viewing the Best Practices Analyzer Results

The event listed in the Best Practices Analyzer pane provides details that will help you to resolve the issue, as shown in Figure 4-27.

Resolution:
In the Windows event log, look for event 12021 for more details about these applications. Consider issuing a new certificate for these applications.

Scan time: 11/11/2014 11:11:00 AM

BPA model version: 2.0.0.0
[More information about this best practice and detailed resolution procedures](#)

Figure 4-27: Viewing Details from the Best Practices Analyzer

The table on TechNet offers the following suggestion for event 12021:

Make sure that the certificate thumbprints that are configured for Web Application Proxy applications are installed on all the Web Application Proxy machines with a private key in the local computer store.

Armed with this information, you can review the certificates on the Web Application Proxy server to ensure that they have the correct names and expiration dates, and that the thumbprint matches the one on the server. Then, you can review the certificates on the server, ensure that they are correct, and reissue them if they are incorrect.

Certificate issues

Certificates play an important role in AD FS and Web Application Proxy. Getting the proper certificates, with the correct names in the certificates on the appropriate machines, is therefore critical to getting Web Application Proxy to function correctly with AD FS.

You might see issues with certificates manifested in error messages like the following:

The trust certificate ("ADFS ProxyTrust - WAP01") is not valid.

There are several possible causes of this issue:

- There might be some sort of network interruption between the Web Application Proxy server and the AD FS server.
- The Web Application Proxy server might have been down for an extended period of time.
- There might be an issue validating the certificate due to problems in the CA infrastructure.
- Time synchronization issues between the Web Application Proxy and AD FS servers might cause them to be out of synchronization.

To resolve these problems, verify the time settings on the Web Application Proxy and AD FS servers and rerun the Install-WebApplicationProxy cmdlets.

Configuration data in AD FS is inconsistent or corrupt

You might also encounter errors for which the configuration data in AD FS could not be found or the data is unusable to the Web Application Proxy server. This can result in errors such as

Configuration data was not found in AD FS.

or

The configuration data stored in AD FS is corrupted or Web Application Proxy was unable to parse it.

or:

Web Application Proxy was unable to retrieve the list of Relying Parties from AD FS.

Several things can cause these errors. It's possible that Web Application Proxy was never fully installed and configured, or there were changes that occurred on the AD FS database that resulted in

corruption. It's also possible that the AD FS server cannot be reached due to a network issue and therefore the AD FS database is not readable.

There are several paths to resolution for these types of errors:

- Run the `Install-WebApplicationProxy` cmdlet again to clear up configuration issues.
- Confirm network connectivity to the AD FS server from the Web Application Proxy server.
- Verify that the `WebApplicationProxy` service is running on the Web Application Proxy server.

Supporting non-SPI-capable clients

Server Name Indication (SNI) is a feature of Secure Sockets Layer (SSL) Transport Layer Security (TLS) that is used in Web Application Proxy server and AD FS to reduce network infrastructure requirements. Traditionally, an SSL certificate had to be bound to an IP address/port combination. This meant that you would need to have a separate IP address configured if you wanted to bind a different certificate to the same port number on a server. With the use of SNI, a certificate is instead bound to the host name and port, allowing you to conserve IP addresses and reduce complexity.

It's important to realize that SNI relies on the requesting client supporting SNI. If the SSL Client Hello doesn't contain the SNI header, `http.sys` won't be able to determine which certificate to offer the client and will reset the connection.

Most modern clients support SNI, but there are some clients that tend to cause issues. Generally, older browsers, legacy operating systems, hardware load balancers, health probes, older versions of WebDAV, ActiveSync on Android, and some older VoIP conferencing devices might be non-SNI-capable devices.

If it is necessary to support non-SNI clients, the easiest solution is to create a fallback certificate binding in `http.sys`. The fallback certificate needs to include any fully qualified domain names (FQDNs) that may need to be supported, including the FQDN for the AD FS service itself (`adfs.contoso.com`), the FQDN of any applications published via Web Application Proxy (`mail.contoso.com`), and the FQDN to support Enterprise registration (`enterpriseregistration.contoso.com`) if you are using Workplace Join.

When you have generated the certificate, get the application GUID and certificate thumbprints in use by using the following Windows PowerShell cmdlet:


```
Get-WebApplicationProxyApplication | fl Name,ExternalURL,ExternalCertificateThumbprint
```

Now that you have the application GUID and certificate thumbprint, you can bind it to the IP wildcard and port 443 by using the following syntax:

```
netsh http add sslcert iport=0.0.0.0:443 certhash=certthumbprint appid={applicationguid}
```

Note that this will need to be run on each server in the AD FS farm, as well as on any Web Application Proxy server.

More info You can find technical details on SNI as a subsection of the TLS Extensions RFC at <https://tools.ietf.org/html/rfc3546#section-3.1>.



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!



Security

For the past several years, cybersecurity has been consistently rated as a top priority for IT. This is not surprising, as top companies and government agencies are being publically called out for being hacked and failing to protect their customers' and employees' personal information.

On the other hand, with readily available tools and lack of adequate protections, attackers are able to infiltrate large organization and remain undetected for a long period of time while conducting exfiltration of secrets or attacking internal resources.

In this chapter, we will explore the layers of protection in Windows Server 2016 Technical Preview that help address emerging threats and make it an active participant in your security defenses. First, we will describe the new shielded virtual machine solution that protects virtual machines (VMs) from attacks on the underlying fabric.

Then, we will introduce you to the extensive threat-resistance components built in to the Windows Server 2016 Technical Preview operating system and the enhanced auditing events that can help security systems detect malicious activity.

Last, we will share with you an end-to-end plan for securing privileged access based on existing and new capabilities in Windows Server.

Shielded VMs

By John Savill

Today, in most virtual environments there are many types of administrators who have access to VM assets, such as storage. That includes virtualization administrators, storage administrators, network administrators, backup administrators, just to name just a few. Many organizations including hosting providers need a way to secure VMs—even from administrators—which is exactly what shielded VMs provides. Note that this protection from administrators is needed for a number of reasons. Here are just a few:

- Phishing attacks
- Stolen administrator credentials
- Insider attacks

Shielded VMs provide protection for the data and state of the VM against inspection, theft, and tampering from those who have administrator privileges. Shielded VM works for Generation 2 VMs that provide the necessary secure startup, UEFI firmware, and virtual Trusted Platform Module (vTPM) 2.0 support required. Although the Hyper-V hosts must be running Windows Server 2016 Technical Preview, the guest operating system in the VM can be Windows Server 2012 or above.

A new Host Guardian Service instance is deployed in the environment, which stores the keys required for an approved Hyper-V host that can prove its health to run shielded VMs.

A shielded VM provides the following benefits:

- BitLocker encrypted drives (utilizing its vTPM)
- A hardened VM worker process (VMWP) that encrypts live migration traffic in addition to its runtime state file, saved state, checkpoints, and even Hyper-V Replica files
- No console access in addition to blocking Windows PowerShell Direct, Guest File Copy Integration Components, and other services that provide possible paths from a user or process with administrative privileges to the VM

How is this security possible? First, it's important that the Hyper-V host has not been compromised before the required keys to access VM resources are released from the Host Guardian Service (HGS). This attestation can happen in one of two ways. The preferred way is by using the TPM 2.0 that is present in the Hyper-V host. Using the TPM, the boot path of the server is assured, which guarantees no malware or root kits are on the server that could compromise the security. The TPM secures communication to and from the HGS attestation service. For hosts that do not have a TPM 2.0, an alternate Active Directory-based attestation is possible; however, this merely checks if the host is part of a configured Active Directory group. Therefore, it does not provide the same levels of assurance and protection from binary meddling and thus host administrator privileges for a sophisticated attacker. However the same shielded VM features are available.

After a host undergoes the attestation, it receives a health certificate from the attestation service on the HGS that authorizes the host to get keys released from the key protection service that also runs on the HGS. The keys are encrypted during transmission and can only be decrypted within a protected enclave that is new to Windows 10 and Windows Server 2016 Technical Preview (more on that later). These keys can then be used to decrypt the vTPM to make it possible for the VM to access its BitLocker-protected storage and start the VM. Therefore, only if a host is authorized and non-compromised will it be able to get the required key and turn on the VM's access to the encrypted storage (not the administrator, though, as the virtual hard drive (VHD) remains encrypted on the drive).

At this point, it might self-defeating: If I am an administrator on the Hyper-V and the keys are released to the host to start the VM, I would be able to gain access to the memory of the host and get the keys, thus nullifying the very security that should protect VMs from administrative privileges. Fortunately, another new feature in Windows 10 and Windows Server 2016 Technical Preview prevents this from happening. This feature is the protected enclave mentioned earlier, which is known as Virtual Secure Mode (VSM). A number of components use this service, including credential guard. VSM is a secure execution environment where secrets and keys are maintained and critical security processes run as Trustlets (small trusted processes) in a secure virtualized partition.

This is not a Hyper-V VM; rather, think of it like a small virtual safe that is protected by virtualization based on technologies such as Second Level Address Translation (SLAT) to prevent people from trying to directly access memory, I/O Memory Management Unit (IOMMU) to protect against Direct Memory Access (DMA) attacks, and so on. The Windows operating system, even the kernel, has no access to VSM. Only white-listed processes (trustlets) that are Microsoft signed are allowed to cross the "bridge" to access VSM. A vTPM Trustlet is used for the vTPM of each VM, separate from the rest of the VM process, which runs in a new type of protected VM worker process. This means that there is no way to

access the memory used to store these keys, even with complete kernel access. If I'm running with a debugger attached, for example, that would be flagged as part of the attestation process, the health check would fail, and the keys would not be released to the host. Remember I mentioned the keys from the key protection service are sent encrypted? It's the VSM where they are decrypted, always keeping the decrypted key protected from the host OS.

When you put all of this together you have the ability to create a secure VM environment that is protected from any level of administrator (when using TPM 2.0 in the host) and will close a security hole many environments cannot close today.

More info To read detailed guides that Microsoft has provided to implement this scenario in your environment, go to <https://gallery.technet.microsoft.com/Shielded-VMs-and-Guarded-44176db3/view/Discussions>.

Threat-resistant technologies

Windows Server 2016 Technical Preview includes integrated threat-resistance technologies that make it an active component in your overall security story. These technologies range from blocking external attackers trying to exploit vulnerabilities (Control Flow Guard) to resistance to attacks by malicious users and software that gained administrator access to the server (Credential Guard and Device Guard). In this section, we dive into some of these new features.

Control Flow Guard

In Windows Server 2016 Technical Preview and Windows 10, the operating system is protected by Control Flow Guard. This highly optimized platform security feature makes it much harder to run arbitrary code through exploits such as *buffer overflows*.

In addition, when a developer compiles his code, the compiler will perform some security checks on the code and then identify the set of functions that are considered a source for an indirect call. These indirect calls might come from a code exploit whereby malformed data is sent into the function, causing it to behave abnormally. The indirect call in non-Control Flow Guard-aware code can cause a memory buffer overrun, which can corrupt other applications or lead to privileged execution. However, because the compiler has identified these sets of functions as potential vulnerabilities and marked them, the runtime will detect and provide additional logic that verifies whether an indirect call is actually valid. If the indirect call validation fails, the application will terminate, preventing the application from causing further damage to the system.

Device Guard (Code Integrity)

In the Windows operating system, there are two modes of operation: kernel mode and user mode. Kernel mode is the one in which the operating system is interacting with the hardware resources on the machine, and user mode is essentially that in which the user experience of running applications takes place. In the past few generations of software that Microsoft has released, it signs its code to ensure that if someone tampers with a binary, it has the ability to detect that something is different. Code Integrity is the part of Windows that performs this function. A common scenario involves hardware drivers. Microsoft requires signed drivers in order to install them onto the operating system so that they can operate in kernel mode. In Windows Server 2016 Technical Preview, further improvements have been made to Code Integrity whereby you can now create policies for your organization's needs. You can deploy these Code Integrity policies to an environment such that an end user cannot download and run untrusted code. This has a direct effective on the spread of malware, given that we know most malware programs typically disguise themselves as something

fun or as a common application. Now, end users will not be able to execute the potentially dangerous code.

You can configure Code Integrity policies individually on machines based on specific needs or you can baseline a machine and capture a *golden image* and use that as the base from which you deploy each additional machine.

You can have only one Code Integrity policy per machine, which is stored in C:\Windows\System32\CodeIntegrity. This means that you might need to implement multiple policies depending on your application requirements. For example, you might have a policy that covers machines in the finance department, a separate one for the engineering department, and so on. The simplest method for most organization's is to create a core application policy based on your software catalog and then where specific additions are needed, you can merge a policy that includes the new software into that specific machine(s).

Code Integrity policies comprise several components, two of which are of particular interest: policy rules and file rules. Policy rules can have a variety of options as detailed in the following table:

Rule option	Description
0 Enabled:UMCI	Code Integrity policies restrict both kernel-mode and user-mode binaries. By default, only kernel-mode binaries are restricted. Turning on this rule option validates user-mode executables and scripts.
1 Enabled:Boot Menu Protection	This option is not currently supported.
2 Required:WHQL	By default, legacy drivers that are not Windows Hardware Quality Labs (WHQL)-signed are allowed to run. Turning on this rule requires that every driver is WHQL-signed and removes legacy driver support. Going forward, every new Windows 10-compatible driver must be WHQL certified.
3 Enabled:Audit Mode (Default)	Allows binaries to run outside of the Code Integrity policy but logs each occurrence in the CodeIntegrity event log, which you can use to update the existing policy before enforcement. To enforce a Code Integrity policy, remove this option.
4 Disabled:Flight Signing	If turned on, Code Integrity policies will not trust flightroot-signed binaries. You would use this would for scenarios in which organizations want to run only released binaries, not flighted builds.
5 Enabled:Inherent Default Policy	This option is not currently supported.
6 Enabled:Unsigned System Integrity Policy (Default)	Allows the policy to remain unsigned. When this option is removed, the policy must be signed and have UpdatePolicySigners added to the policy to make future policy modifications possible.
7 Allowed:Debug Policy Augmented	This option is not currently supported.
8 Required:EV Signers	In addition to requiring that drivers be WHQL-signed, this rule requires that they must have been submitted by a partner that has an Extended Verification (EV) certificate. All future Windows 10 and later drivers will meet this requirement.
9 Enabled:Advanced Boot Options Menu	The F8 prestartup menu is turned off by default for

	all Code Integrity policies. Setting this rule option allows the F8 menu to appear to physically present users.
10 Enabled:Boot Audit on Failure	Used when the Code Integrity policy is in enforcement mode. When a driver fails during startup, the Code Integrity policy will be placed in audit mode so that Windows will load. Administrators can validate the reason for the failure in the CodeIntegrity event log.

With file rules, you can configure the level at which you want to trust the application. File rules are specified when you create a new Code Integrity policy from a scan and when you create a policy from audit events. If you have a base Code Integrity policy in place already and you want to merge additional policies, the files will be combined, as well. Use the following table to help you select the level of trust for an application that you want to configure.

Rule level	Description
Hash	Specifies individual hash values for each discovered binary. Although this level is specific, it can cause additional administrative overhead to maintain the current product versions' hash values. Each time a binary is updated, the hash value changes, therefore requiring a policy update.
FileName	Specifies individual binary file names. Although the hash values for an application are modified when updated, the file names typically are not. This offers less specific security than the hash level but does not typically require a policy update when any binary is modified.
SignedVersion	This combines the publisher rule with a file version number. This option allows anything from the specified publisher, with a file version at or above the specified version number, to run.
Publisher	This is a combination of the PCA certificate and the common name (CN) on the leaf certificate. In the scenario that a PCA certificate is used to sign multiple companies' applications (such as VeriSign), this rule level allows organizations to trust the PCA certificate but only for the company whose name is on the leaf certificate (for example, Intel for device drivers). This level trusts a certificate with a long validity period but only when combined with a trusted leaf certificate.
FilePublisher	This is a combination of the publisher file-rule level and the SignedVersion rule level. Any signed file from the trusted publisher that is the specified version or newer is trusted.
LeafCertificate	Adds trusted signers at the individual signing certificate level. The benefit of using this level versus the individual hash level is that new versions of the product will have different hash values but typically the same signing certificate. Using this level, no policy update would be needed to run the new version of the application. However, leaf certificates have much shorter validity periods than PCA certificates, so additional administrative overhead is associated with updating the code integrity policy when these certificates expire.
PcaCertificate	Adds the highest certificate in the provided certificate chain to signers. This is typically one certificate below the root certificate, because the scan does not

	validate anything above the presented signature by going online or checking local root stores.
RootCertificate	Currently unsupported.
WHQL	Trusts binaries if they have been validated and signed by WHQL. This is primarily for kernel binaries.
WHQLPublisher	This is a combination of the WHQL and the CN on the leaf certificate and is primarily for kernel binaries.
WHQLFilePublisher	Specifies that the binaries are validated and signed by WHQL, with a specific publisher (WHQLPublisher), and that the binary is the specified version or newer. This is primarily for kernel binaries.

Building a golden image

Now, let's look at an example of building a golden image policy from a reference machine. It is important to note that you need to ensure that this is a freshly built machine, free of malware and viruses.

From an elevated Windows PowerShell prompt, type the following information to begin preparing for the initial scan:

```
$CIPolicyPath= "C:\Temp"
$InitialCIPolicy=$CIPolicyPath+"InitialScan.xml"
$CIPolicyBin=$CIPolicyPath+"DeviceGuardPolicy.bin"
```

These commands set up to store the output of a scan we will initialize in this next code example:

```
New-CIPolicy -Level PcaCertificate -FilePath $InitialCIPolicy -UserPEs 3> CIPolicyLog.txt
```

The UserPEs option automatically turns on User Mode Code Integrity; this will scan the machine on the level PcaCertificate. New-CIPolicy has various other parameters, which you can find at <https://technet.microsoft.com/library/mt634473.aspx>.

The next step is to convert this into a binary format for later use. The follow code performs this step:

```
ConvertFrom-CIPolicy $InitialCIPolicy $CIPolicyBin
```

If you check on the temp drive where we specified, you will see the .bin file and the .xml file from the scan. Save these to a secure location for later use.

More info For some basic information on how to get started with Code Integrity policies as well as further information about creating an audit policy and deploying it via Group Policy, go to [https://technet.microsoft.com/library/mt463091\(v=vs.85\).aspx#code_integrity_policies](https://technet.microsoft.com/library/mt463091(v=vs.85).aspx#code_integrity_policies).

Credential Guard

Credential Guard isolates secrets using virtualization-based technologies so that only privileged systems can access them. Credential Guard offers the following features:

- **Hardware security** This increases the security of derived domain credentials by taking advantage of platform security features, including, Secure Boot and virtualization.
- **Virtualization-based security** Windows services that manage derived domain credentials and other secrets run in a protected environment that is isolated from the running operating system.

- **Better protection against advanced persistent threats** Secures derived domain credentials by using the virtualization-based security. This blocks the credential theft attack techniques and tools used in many attacks. Malware running in the operating system with administrative privileges cannot extract secrets that are protected by virtualization-based security.
- **Manageability** Manage by using Group Policy, WMI, from a command prompt, and Windows PowerShell.

Normally, secrets are stored in the memory of the Local Security Authority (LSA) process in Windows. With Credential Guard, the LSA talks to a new component called *isolated LSA*. This isolated LSA is virtualization-based and is not accessible by the rest of the operating system. In Figure 5-1 we show you the isolation provided by the Virtualization-based security for the LSASS process with respect to the LSASS process.




Figure 5-1: Virtualization-based LSA process

When Credential Guard is turned on, older variants of NTLM or Kerberos (i.e., NTLM v1, MS-CHAPv2, etc.) are no longer supported.

Because Credential Guard is virtualization-based, it requires some specific hardware support. The following table details some of those requirements:

Requirement	Description
UEFI firmware version 2.3.1 or higher and Secure Boot	To verify that the firmware is using UEFI version 2.3.1 or higher and Secure Boot, you can validate it against the System.Fundamentals.Firmware.CS.UEFI Secure Boot.Connected Standby Windows Hardware Compatibility Program requirement.
Virtualization extensions	The following virtualization extensions are required to support virtualization-based security: • Intel VT-x or AMD-V • Second Level Address Translation
x64 architecture	The features that virtualization-based security uses in the Windows hypervisor can run only on a 64-bit PC.
A VT-d or AMD-Vi IOMMU	In Windows 10, an IOMMU enhances system resiliency against memory attacks. TPM 1.2 and 2.0 provides protection for encryption keys that are stored in the firmware. TPM 1.2 is not supported on Windows 10 (Build 10240); however, it is supported in Windows 10, Version 1511 (Build 10586) and later.
Trusted Platform Module (TPM) version 1.2 or 2.0	Note: If you don't have a TPM installed, Credential Guard will still be turned on, but the keys used to encrypt Credential Guard will not be protected by the TPM.

Secure firmware update process	To verify that the firmware complies with the secure firmware update process, you can validate it against the System.Fundamentals.Firmware.UEFI SecureBoot Windows Hardware Compatibility Program requirement.
The firmware is updated for Secure MOR implementation	Credential Guard requires the secure MOR bit to help prevent certain memory attacks.
Physical PC	For PCs running Windows 10, you cannot run Credential Guard on a VM.

The simplest way to get Credential Guard implemented for your organization is to turn it on via Group Policy and designate the machines in your enterprise for which you want to apply it.

From the Group Policy Management Console, create a new group policy or edit an existing one. Then, go to Computer Configuration > Administrative Templates > System > Device Guard.

Double-click Turn On Virtualization Based Security, and then, in the dialog box that opens (see Figure 5-2), select the Enabled option. In the Select Platform Security Level list box, choose Secure Boot or Secure Boot And DMA Protection. In the Credential Guard Configuration list box, select Enabled With UEFI lock, and then click OK. If you want to be able to turn off Credential Guard remotely, choose Enabled Without Lock from the Credential Guard Configuration list box instead of Enabled With UEFI lock.




Figure 5-2: Group Policy options for Credential Guard

More info For further information, go to [https://technet.microsoft.com/library/mt483740\(v=vs.85\).aspx](https://technet.microsoft.com/library/mt483740(v=vs.85).aspx).

Windows Defender

Windows Defender is included (and running) by default when you install Windows Server 2016 Technical Preview. Depending on the SKU of Windows Server 2016 Technical Preview that you choose to install, it might or might not have the GUI tools. If it doesn't have the GUI tools, you can install them via Windows PowerShell, as follows:

```
Install-WindowsFeature -Name Windows-Defender-GUI
```

If your organization has a company standard for malware technology, you can uninstall it by using Windows PowerShell, as well:

```
Uninstall-WindowsFeature -Name Windows-Server-Antimalware
```

Windows Defender receives updates via Windows Update. If your organization manages Windows Update via an update deployment tool, you need to ensure that you are downloading the updates to keep Windows Defender up to date with its definitions.

You also can configure Windows Defender via Group Policy for central control and administration.

Threat detection technologies

No matter how much you try to secure an environment, you still need to perform audits to validate whether those measures are effective. In Windows Server 2016 Technical Preview, two new audit subcategories have been added to give greater insight into the events:

- **Audit Group Membership** This is part of the Logon/Logoff event category. The events in this subcategory are generated when group memberships are enumerated or queried on the PC where the sign-in session was created.
- **Audit PNP Activity** Found in the Detailed Tracking category, you can use the Audit PNP Activity subcategory to audit when plug-and-play detects an external device. Only Success audits are recorded for this category.

Additional changes have been made in Windows Server 2016 Technical Preview that expose more information to help you identify and address threats quickly. The following table provides more information:

Area	Improvements
Kernel Default Audit Policy	In previous releases, the kernel depended on the LSA to retrieve information in some of its events. In Windows 10, the process creation events audit policy is automatically turned on until an actual audit policy is received from LSA. This results in better auditing of services that might start before LSA starts
Default process Security ACL (SACL) to LSASS.exe	A default process SACL was added to LSASS.exe to log processes attempting to access LSASS.exe. The SACL is L"S:(AU;SAFA;0x0010;;;WD)". You can turn this on under Advanced Audit Policy Configuration\Object Access\Audit Kernel Object.
New fields in the sign-in event	The sign-in event ID 4624 has been updated to include more verbose information to make them easier to analyze. The following fields have been added to event 4624: • MachineLogon String: yes or no If the account that signed in to the PC is a computer account, this field will be yes; otherwise, the field is no.

	• ElevatedToken String: yes or no If the account that signed in to the PC is an administrative sign-in, this field will be yes; otherwise, the field is no. Additionally, if this is part of a split token, the linked login ID (LSAP_LOGON_SESSION) will also be shown. • TargetOutboundUserName String and TargetOutboundUserDomain String The user name and domain of the identity that was created by the LogonUser method for outbound traffic. • VirtualAccount String: yes or no If the account that signed in to the PC is a virtual account, this field will be yes; otherwise, the field is no. • GroupMembership String A list of all of the groups in the user's token. • RestrictedAdminMode String: yes or no If the user signs in to the PC in restricted admin mode with Remote Desktop, this field will be yes.
New fields in the process creation event	The sign-in event ID 4688 has been updated to include more verbose information to make it easier to analyze. The following fields have been added to event 4688: • TargetUserSid String The SID of the target principal. • TargetUserName String The account name of the target user. • TargetDomainName String The domain of the target user. • TargetLogonId String The logon ID of the target user. • ParentProcessName String The name of the creator process. • ParentProcessId String A pointer to the actual parent process if it's different from the creator process.
Security Account Manager (SAM) events	New SAM events were added to cover SAM APIs that perform read/query operations. In previous versions of Windows, only write operations were audited. The new events are event ID 4798 and event ID 4799. The following APIs are now audited: - SamrEnumerateGroupsInDomain - SamrEnumerateUsersInDomain - SamrEnumerateAliasesInDomain - SamrGetAliasMembership

	SamrLookupNamesInDomain SamrLookupIdsInDomain SamrQueryInformationUser SamrQueryInformationGroup SamrQueryInformationUserAlias SamrGetMembersInGroup SamrGetMembersInAlias SamrGetUserDomainPasswordInformation
Boot Configuration Database (BCD) events	Event ID 4826 has been added to track the following changes to the BCD: DEP/NEX settings Test signing PCAT SB simulation Debug Boot debug Integrity Services Disable Winload debugging menu
PNP Events	Event ID 6416 has been added to track when an external device is detected through plug-and-play. One important scenario is if an external device that contains malware is inserted into a high-value machine that doesn't expect this type of action, such as a domain controller.

Securing privileged access

In this section, we are going to explore a few concepts regarding securing privileged access. First we are going to dive into the concepts of Just In Time and Just Enough Administration (JEA). Then, we are going to explain how you combine all of the tools and technologies we have discussed in this chapter into an implementation strategy for your organization.

Just In Time and Just Enough Administration

Just In Time (JIT) administration is a fairly basic concept: the principal is that we evolve to a state in which there are no full-time administrators, or more specifically we have no accounts that have full-time administrator privileges. Rather, through a simple process, the privileges required are requested just before they are actually needed, then approved, and then granted to the account for a specific time period. This ensures that the task can be carried out successfully with the correct amount of privileges for the allotted time. JIT works in conjunction with Just Enough Administration (JEA) to secure the correct privileges. In Windows Server 2016 these technologies are combined to provide Privileged Access Management (PAM).

More info For more information about PAM, go to <https://technet.microsoft.com/library/dn903243.aspx>.

Now, let's take a quick look at JEA. This is part of the Windows Management Framework 5.0 package and has been supported since Windows Server 2008 R2. Using JEA, you can assign specific privileges (just enough of them) to a user account to perform a given required function. This means that you don't need to assign a user to an administrator account and then remember to remove them later. JEA gives us the role-based access control (RBAC) that modern enterprises require to achieve more secure environments.

To implement JEA on a system, you first need a Windows PowerShell Session Configuration file. Use the `New-PSSessionConfigurationFile` cmdlet to create the .pssc file you need to control access by running the following syntax:

```
New-PSSessionConfigurationFile -Path "$env:Programdata\<JEAConfigDirectory>\<filename>.pssc"
```

The following is a sample of the default configuration file this command generates:

```
@{
# Version number of the schema used for this document
SchemaVersion = '2.0.0.0'
# ID used to uniquely identify this document
GUID = '1da190ce-fc94-4f8b-98e0-7d70fd9154b1'
# Author of this document
Author = 'john'
# Description of the functionality provided by these settings
# Description = ''

# Session type defaults to apply for this session configuration. Can be 'RestrictedRemoteServer'
(recommended), 'Empty', or 'Default'
SessionType = 'Default'
# Directory to place session transcripts for this session configuration
# TranscriptDirectory = 'C:\Transcripts\'
# Whether to run this session configuration as the machine's (virtual) administrator account
# RunAsVirtualAccount = $true
# Groups associated with machine's (virtual) administrator account
# RunAsVirtualAccountGroups = 'Remote Desktop Users', 'Remote Management Users'
# Scripts to run when applied to a session
# ScriptsToProcess = 'C:\ConfigData\InitScript1.ps1', 'C:\ConfigData\InitScript2.ps1'
# User roles (security groups), and the role capabilities that should be applied to them when applied to a
session
# RoleDefinitions = @{ 'CONTOSO\SqlAdmins' = @{ RoleCapabilities = 'SqlAdministration' } };
'CONTOSO\ServerMonitors' = @{ VisibleCmdlets = 'Get-Process' } }
}
```

The core areas of interest to change are the `SessionType`, which is set to `Default`. For JEA to work, you need to configure this as `RestrictedRemoteServer`. Next, you need to uncomment `# RunAsVirtualAccount = $True`, which ensures that the session will have “virtual” administrator privileges. Finally, you need to modify the `RoleDefinitions` section and uncomment it to reflect your environment.

After you generate and set up the configuration file, you need to register it by using the `Register-PSSessionConfiguration` cmdlet.

```
Register-PSSessionConfiguration -Name <Name> -Path "$env:Programdata\<JEAConfigDirectory>\<filename>.pssc"
```

You can test this by connecting to the machine as you would a normal Windows PowerShell remote session.

```
Enter-PSSession -ComputerName <ComputerName> -ConfigurationName <JEAConfigName> -Credential $cred
```

You can create another file called a Role Capability file with the extension .psrc. You can use this file to define what commands and applications are visible to the specific roles you define. You use the `New-PSRoleCapabilityFile` cmdlet to create a blank template.

This file contains sections in which you can define which modules to import and which functions and cmdlets are exposed.

```
# ModulesToImport = 'MyCustomModule',
@{ ModuleName = 'MyCustomModule2'; ModuleVersion = '1.0.0.0';
GUID = '4d30d5f0-cb16-4898-812d-f20a6c596bdf'
}
# VisibleFunctions = 'Invoke-Function1',
@{ Name = 'Invoke-Function2';
Parameters = @{ Name = 'Parameter1';
ValidateSet = 'Item1', 'Item2' },
@{ Name = 'Parameter2'; ValidatePattern = 'L*' } }
# VisibleCmdlets = 'Invoke-Cmdlet1',
@{ Name = 'Invoke-Cmdlet2'; Parameters = @{ Name = 'Parameter1'; ValidateSet = 'Item1', 'Item2' },
@{ Name = 'Parameter2'; ValidatePattern = 'L*' } }
```

More info JEA is a detailed subject, and we can provide only the basics here. For further guidance and to see all of the configuration options, go to <http://aka.ms/JEA>.

A strategy for securing privileged access

It has to be said that no matter how secure you can make an operating system or service it is only as secure as the weakest password. For example, suppose that you have the most sensitive data on earth and you encrypt it by using the most sophisticated technology, but then you use a password like "Password01"; this utterly defeats the purpose of putting in place a battery of secure technologies.

Let's look at another scenario. Walk around your office and count how many people have written their passwords on notes and stuck them on their keyboards or monitors. Then, observe how many people have pictures of their family or pets on their desk. When those people need to think of a password, what is the likelihood that it might be something personal based on the pictures?

Now, let's consider a final scenario: the social engineering attack. With this particular form of attack—which is a leading cause of security breaks—the attacker calls someone, out of the blue, and pretends to be from IT, saying he needs to verify some account information. If the attacker is good at his job, the chances are high that the hapless victim will readily provide the information.

With those scenarios in mind, the attacker will gain access to something and potentially use that access that to perform an escalated attack. But, what if the account were a privileged one in the first place.

Securing privileged access is not a single technology; it is a set of practices that an organization can implement to become more secure. Although focused primarily on privileged access, it highlights the need for any organization to implement and test all policies related to security and conduct the necessary readiness to make people aware of potential areas of exposure.

No network to which users have access will ever be 100 percent secure, but to begin down the path of securing privileged access to systems and networks, you must be diligent with regard to the following basics:

- **Updates** Deploy updates to domain controllers within seven days of release.
- **Remove users as local administrators** Monitor and remove users from local administrators if they don't need this access. Use Active Directory to control membership centrally, if required.
- **Baseline security policies** Deploy policies that will maintain a standard configuration for the organization. Exceptions will exist, of course, based on applications and certain requirements, but these should be challenged on a repeated basis to ensure the system is as compliant as possible.
- **Antimalware programs** Maintain regular updating and regular scans of the environment. Clean and remove threats as quickly as possible.
- **Log and analysis** Capture security information, perform regular reviews, and identify anomalies within the log set. Perform follow-up action on each detected item to ensure that it is an identified source and safe "risk."
- **Software inventory and deployment** Controlling the software installed in an environment is paramount to ensure that end users don't install malware into the environment. In the same manner, it is important to know what software is out there and maintain an inventory so that you are aware if the state of a system has changed.

With these basics covered, we can move into more details about the strategy that underpins securing privileged access. Be aware that you will not achieve this strategy overnight, and this should be built as a progressive implementation so that the organization's practices can change and adapt to these new principles.

As with most strategies, you need to establish short-, medium- and long-term goals. The following table describes the goals and the time frames you should use, and the areas of focus for each goal.

Goal	Time frame	Description
Short term	2- to 4-week plan	Quick mitigation of the most frequently used attacks
Medium term	1- to 3-month plan	Build visibility and control of administrative activity
Long term	6 months and beyond	Build a proactive security posture

Short-term plan

For the short-term goal, it is critical that you mitigate the most frequently used attacks in any organization to provide a secure base.

One of the first things you need to do is to establish *separation of duties*. This means that if you need to perform a privileged-access task, you should have an appropriate privileged-access account to carry it out. You should never grant your standard user account privileged access in a network to perform tasks. This account should always be considered a user. The privileged-access account you create for tasks can be audited and tracked in more detail. Because you maintain a different set of credentials for this account with stricter requirements, you will be able to mitigate an attack if your user account is compromised.

Securing the local administrator account was previously done during deployment and was rarely changed after it was set. The password was usually kept the same throughout the entire estate of workstations, which led to a huge problem if the password was compromised. However, if you don't use the same password throughout the estate, you might have a more complicated problem trying to remember the unique password for each of the workstations. To help you manage the local administrator password for both workstations and servers, Microsoft provides a tool called Local Administrator Password Solution (LAPS).

LAPS creates a unique password for each server and workstation in an environment and stores them in Active Directory as a confidential attribute in the computer object. They have an appropriate access control lists applied to them so that only the appropriate accounts can access them and retrieve them as necessary. For more information on LAPS, go to <http://aka.ms/LAPS>.

The final key part of the short-term goals should be focused around creating privileged access workstations (PAWs). PAWs are hardened workstations implemented specifically to act as a controlled point of administration to more secure systems. PAWs would be restricted from accessing the Internet or unsecure resources ensuring that their attack surface is to an absolute minimum. Only a restricted set of authorized users would also be able to sign in to the PAWs, which in turn would reduce the ability to attack secure part of the networks. For more information on PAWs, go to <http://aka.ms/CyberPAW>.

Figure 5-3 illustrates the steps that you can take as part of your short-term plan.




Figure 5-3: Short-term goal plan

The figure shows four separate areas:

1. Create a separate administrator account for administrative tasks, as shown with the Admin and User.
2. Deploy PAWs for Active Directory administrators. For more information, go to <http://aka.ms/cyberPAW>, where this step is shown as Phase 1.
3. Create unique LAPS for workstations. For more information, go to <http://aka.ms/LAPS>.
4. Create unique LAPS for Servers. For more information, go to <http://aka.ms/LAPS>.

Medium-term plan

The first thing you need to do for your medium-term plan is to expand the deployment of PAWs so that you can bring more systems into scope, which you can manage only from these workstations. Following on from that, you should begin to focus on implementing time-bound privileges; that is, a user can request privileges that will expire after a predefined period of time. This means there does not need to be actual administrators, as such, because the users can request the access they need, be approved, and perform the necessary tasks. This concept is based on Microsoft Identity Manager and functions provided by JEA.

You also should implement multifactor authentication for privileged access to further mitigate attacks on the systems. You can do this by using token-based security or call-back or smart cards. Next, you can begin to implement JEA. JEA is simple in principle because it specifies that you grant the very minimal amount of privileges to an account that are needed to perform the given function.

The next step is to further secure domain controllers, and you will finish by implementing threat detection via Advanced Threat Analytics (ATA). ATA provides the ability to detect abnormal behavior in your systems and make you aware of them quickly. It does this by profiling your user's behavior and establishing what that user's normal patterns are. If the user does something outside this normal pattern, ATA will alert you. ATA is far more advanced than this simple explanation implies. To learn more about it, go to <http://aka.ms/ata>.

Figure 5-4 presents an illustrated overview of the medium-term plan.




Figure 5-4: Medium-term goal plan

The figure shows six separate areas:

1. Extend PAWs to all administrators and provide additional hardening such as Credential Guard and RDP Restricted Admin. For more information, go to <http://aka.ms/CyberPAW>, where this is shown in Phases 2 and 3.
2. Establish time-bound privileges (no permanent administrators). For more information, go to <http://aka.ms/AzurePIM>.
3. Create multifactor elevation. For more information, go to <http://aka.ms/PAM>.
4. Provide JEA for domain controller maintenance. For more information, go to <http://aka.ms/JEA>.
5. Lower the attack surface of domains and domain controllers. For more information, go to <http://aka.ms/HardenAD>.
6. Implement Attack Detection for your servers and domain controllers. For more information, go to <http://aka.ms/ata>.

Long-term plan

The long-term goals (see Figure 5-5) detail the final parts to date in an ever-evolving strategy. Securing your environment never stops. Therefore, this strategy will need to be reviewed and adapted over time, but it will provide you with a basis to begin and grow.

As with software development, you should apply a lifecycle with regard to how you control access to resources. Your approach should be based the latest principles and JEA. Following on from this, all administrators should be issued strong authentication mechanisms such as SmartCard or Passport Authentication.

To really enhance protection, you can implement a secure forest that is isolated from a traditional user forest. Here, you can store the most secure systems in the environment and be fully isolated from the production network. The next section is to implement code integrity, which will ensure that only authorized code can be run on the systems.

Finally, you can use Shielded VMs. In this case, you can begin by focusing on domain controllers so that an attacker can't inspect a VM and copy it from the drives, or carry out a host attack to gain access to the VM.




Figure 5-5: Long-term goal plan

The figure identifies the following areas:

1. Modernize roles and the delegation model
2. Implement smart card or passport authentication for all administrators (<http://aka.ms/passport>)
3. Create a specific administrator forest for Active Directory administrators (<http://aka.ms/ESAE>)
4. Implement a code-integrity policy for domain controllers in Windows Server 2016 Technical Preview
5. Implement Shielded VMs for domain controllers in Windows Server 2016 Technical Preview and Hyper-V Fabric (<http://aka.ms/shieldedvms>)

App Plat

In Windows Server 2016 Technical Preview, Microsoft takes a focused approach when it comes to offering an App Platform for our customers. With the introduction of two new technologies, *Nano Server* and *containers*, you can now take advantage of a highly optimized, scalable, and secure experience for App Plat technologies.

Nano Server

By Andrew Mason

Nano Server is an exciting new installation option for Windows Server 2016 Technical Preview that has an even smaller footprint than the Server Core installation option. In this section, I explain what Nano Server is, why it was created, and how you can deploy and manage it.

Understanding Nano Server

Nano Server is a new, small-footprint, headless installation option for Windows Server 2016. It is a deep refactoring of Windows Server that is optimized for the cloud. As such, Nano Server in Windows Server 2016 Technical Preview is focused on two primary scenarios: the Microsoft cloud platform system (compute and storage hosts) and the platform for born-in-the-cloud applications. Because Nano Server is focused on these two scenarios and is fully headless, it might require some changes to management and operations procedures for organizations that aren't fully managing their current server deployments remotely.

Windows Server customers have provided this feedback:

- Restarts have a negative impact on my business—why do I need to restart because of a patch to a feature I never use?
- When a restart is required, my servers need to be back in service as soon as possible.
- Large server images take a long time to deploy and consume a lot of network bandwidth.
- If the operating system consumes fewer resources, I can increase my virtual machine density.
- We can no longer afford the security risks of the “install everything everywhere” approach.

Nano Server addresses these problems by including just the functionality required for its proposed use cases and nothing more. This minimizes the patch surface area, thus eliminating restarts and minimizing the footprint, which provides faster deployment and restart time and frees up resources for other uses.

Reduced restarts

To support the need for fewer restarts, the most frequently serviced binaries were removed from Nano Server, because servicing is the most common cause of restarts. Analyzing the patches released in 2014 for Windows Server highlights the differences between Nano Server, Server Core, and Full Server installation options. As Figure 6-1 shows Nano Server provides significant improvement in the number of applicable patches, dramatically reducing restarts.




Figure 6-1: Patch comparison between Nano Server, Server Core, and Full Server

Security improvements

In addition to fewer patches, eliminating installed-by-default functionality from Nano Server also reduces the number of drivers, services, and ports open, as shown Figure 6-2.




Figure 6-2: Default functionality comparison between Nano Server and Server Core

Resource utilization

Minimizing the resources utilized by Nano Server frees resources that can be used to increase virtual machine (VM) density. It also improves startup performance when restarts are required, as demonstrated in Figure 6-3.




Figure 6-3: Resource utilization comparison between Nano Server and Server Core

Deployment improvements

Setup time, including specialization, for Nano Server is significantly less than for Server Core, as is the footprint (see Figure 6-4). This provides fast deployments with less to copy over the network when redeploying, reducing the network bandwidth for deployments that must be accounted for in overall capacity.




Figure 6-4: Deployment requirements comparison between Nano Server and Server Core

Server Core pattern

Figure 6-5 shows how Nano Server will be a separate installation option from the other server installation options in Windows Server 2016 Technical Preview, much as Server Core was in Windows Server 2008 and Windows Server 2008 R2.




Figure 6-5: Architecture of different installation options for Windows Server 2016 Technical Preview

Note Moving between Nano Server and the other installation options will require reinstallation.

Deploying Nano Server

Nano Server is a new installation option for Windows Server 2016 Technical Preview; however, unlike Server Core, it does not appear as an option when you run setup. This is because Nano Server requires that you customize the image for your hardware and the role it will play before deploying, as is discussed further later in the chapter. You can find Nano Server on the Windows Server 2016 Technical Preview media, in the \Nano Server folder; all of the packages that you can install on Nano Server are in \Nano Server\Packages.

More info For the latest information regarding Nano Server deployment, see the Nano Server guide at <https://msdn.microsoft.com/library/mt126167.aspx>.

Drivers

Because Nano Server does not have user-mode plug-and-play, it is necessary for you to add the drivers for your hardware to the Nano Server image before you deploy it. Nano Server uses the same drivers as Windows Server, so you can use any supported hardware that has a driver for Windows Server with Nano Server, including the following:

- Network adapters
- Storage controllers
- Drives

Although there is no need for a special Nano Server version of the driver, if the hardware requires a special tool for configuration and the current tool does not work remotely, the hardware vendor will need to provide an updated tool or instructions for configuration on Nano Server.

Note You add drivers to the Nano Server image by using Dism /add-driver.

Roles and features

Nano Server also separates the package store from the image. Therefore, none of the role or feature binaries are in the WinSXS folder when Nano Server is deployed; you must add them to the image prior to deploying Nano Server. This makes it possible for you to configure the deployed Nano Server image with only what is necessary for the role of the server. To configure Nano Server to be a cloud platform system compute or storage host, the following are available for inclusion in the image:

- Hyper-V
- Scale-out file server
- Clustering
- IIS
- DNS
- Containers

Note You install roles and features by using Dism /add-package.

Additional roles will be added over time. You can check to validate what are the latest roles added to Nano Server Support at <https://msdn.microsoft.com/library/mt126167.aspx>.

Specializing Nano Server

Just like Server Core, you can use a subset of what is available via Unattend to specialize a Nano Server image. In an effort to further reduce the deployment time beyond just the smaller footprint, a couple of commonly used unattend settings are available to set offline:

- Computer name
- Domain join using Djoin.exe

More info For information on how to perform offline domain join using Djoin.exe, see [https://technet.microsoft.com/library/offline-domain-join-djoin-step-by-step\(v=ws.10\).aspx](https://technet.microsoft.com/library/offline-domain-join-djoin-step-by-step(v=ws.10).aspx).

When you deploy a Nano Server image with these settings configured in the offline section of the unattend file, Nano Server is specialized on first startup. This eliminates the second startup that occurs with Server Core during specialization, further reducing deployment time.

Remotely managing Nano Server

Nano Server is truly headless—there is no way to sign in locally or to use Remote Desktop to connect remotely. Both of these have dependencies that would require including frequently serviced features, so they are not available in Nano Server. As a result, you must perform all Nano Server management remotely, either via Windows PowerShell, Windows Management Instrumentation (WMI), Windows Remote Shell (WinRS), Emergency Management Services (EMS), or remote GUI tools.

Note The options discussed in this section to remotely manage Nano Server are the same mechanisms that you can use to remotely manage Server Core. The only difference in managing Nano Server is that you must do it all remotely.

Windows PowerShell

Nano Server includes a refactored subset of Windows PowerShell called Core PowerShell, which is based on the CoreCLR. Core PowerShell provides the following:

- Full Windows PowerShell language compatibility
- Full Windows PowerShell remoting
- Most core engine features
- Support for all cmdlet types, including C#, Windows PowerShell, and CIM

Because Nano Server includes Core PowerShell, it is possible to use PowerShell Remoting to manage Nano Server. To do so, you need to be an administrator on the Nano Server machine and add its IP address to the management machine's trusted hosts. To do that, from an elevated Windows PowerShell prompt, run the following command (which, for this example, assumes the Nano Server machine's IP address is 192.168.1.10):

```
PS C:\> Set-Item WSMan:\localhost\Client\TrustedHosts "192.168.1.10"
```

Next, you need to make yourself an administrator on the Nano Server machine by using the following WinRS command (chcp is used to set the code page) from an elevated command prompt:

```
C:\> chcp 65001
C:\> Winrs -r:192.168.1.10 -u:192.168.1.10\Administrator -p:NanoServer net localgroup Administrators /add
domain\username
```

To use CredSSP, you need to turn on CredSSP on the Nano Server machine (it's off by default, for security reasons):

```
C:\> winrm s winrm/config/service/auth @{CredSSP="true"} -r:192.168.1.10 -u:domain\username
```

The following is an example of how to start an interactive remoting session:

```
PS C:\NanoServer> $ip = "192.168.1.10"
PS C:\NanoServer> $user = "Administrator"
PS C:\NanoServer> Enter-PSSession -ComputerName $ip -Credential $user
```

After you have done this, you can now run any available Windows PowerShell command as if you were entering it directly into the Nano Server console; for example:

```
[192.168.1.10]: PS C:\users\user1\Documents> Get-Process w*
[192.168.1.10]: PS C:\users\user1\Documents> ipconfig /all
```

Not all Windows PowerShell commands are available in Nano Server. To see which cmdlets are available, run the following command:

```
[192.168.1.10]: PS C:\users\user1\Documents> Get-Command -CommandType Cmdlet
```

To end the remoting session, run this command:

```
[192.168.1.10]: PS C:\users\user1\Documents> Exit-PSSession
```

WMI

Nano Server includes WMI v1 and WMI v2 as well as the providers for the included functionality.

WMIC

You also can use Wmic.exe to run WMI commands remotely. Note that you must have a non-null administrator password, as in the following example:

```
C:\NanoServer> wmic /user:Administrator /password:Tuva /Node:192.168.1.10 OS get Name
```

More info You can read more about Wmic.exe at <http://technet.microsoft.com/library/cc754534.aspx>.

WinRM

You can use CIM sessions and CIM instances in Windows PowerShell to run WMI commands remotely over WinRM, as demonstrated here:

```
PS C:\NanoServer> $ip = "192.168.1.10"
PS C:\NanoServer> $user = "Administrator"
PS C:\NanoServer> $pwd = "Tuva"
PS C:\NanoServer> $cim = New-CimSession -Credential $user -ComputerName $ip
```

When the CIM session is established, you can run various WMI commands, such as the following:

```
PS C:\NanoServer> Get-CimInstance -CimSession $cim -ClassName Win32_ComputerSystem | Format-List *
PS C:\NanoServer> Get-CimInstance -Query "SELECT * from Win32_Process WHERE name LIKE 'p%'"
```

WinRS

Using Windows Remote Management (WinRM), you can run programs remotely. Before you can use WinRS, you need to configure the WinRM service and set the code page, as follows:

```
C:\NanoServer> winrm quickconfig
C:\NanoServer> winrm set winrm/config/client @{TrustedHosts="*"}
C:\NanoServer> chcp 65001
```

After you configure the WinRM service, you can run commands remotely, as if you were running them from the command line:

```
C:\NanoServer> winrs -r:192.168.1.10 -u:Administrator -p:Tuva ipconfig
```

More info To learn more about about WinRS, go to <http://technet.microsoft.com/library/hh875630.aspx>.

EMS

EMS is yet another tool that you can use to remotely manage Nano Server by connecting a serial cable between the management machine and the Nano Server machine.

After you set up EMS in the Boot Configuration Data (BCD) settings for the Nano Server machine and start Nano Server, start a terminal emulator, such as PuTTY, from an elevated prompt on the management machine, and then complete the following steps:

1. Set the speed to the same baud rate you used in the Nano Server BCD.
2. Select Serial for Connection Type.
3. Provide the correct value for Serial Line.

Remote GUI tools

In addition to the command-line remote management options discussed in the previous sections, you can use many existing remote GUI tools to remotely manage Nano Server. Because there is no local sign-in or Remote Desktop in Nano Server and there are tools that even in Server Core don't have remote GUI replacements, for example Task Manager, there are a set of web-based remote GUI replacements under development. You can use these web-based remote GUI tools to manage Nano Server as well as Server Core and any of the other installation options.

Containers

By John McCabe

This section introduces a new technology to Windows Server 2016 Technical Preview called *containers*. Containers come in two types:

- Windows Server containers
- Hyper-V containers

In this section we explain what a container is and why is it important?

What is a container?

A container in its simplest form is exactly that—a container. It is an isolated environment in which you can run an application without fear of changes due to applications or configuration. Containers share key components (kernel, system drivers, and so on) that can reduce startup time and provide greater density than you can achieve with a VM. Figure 6-6 shows an abstract of a container.




Figure 6-6: Conceptual container layout

As the illustration demonstrates, a host operating system can host many containers and allow them to be completely isolated while sharing key components of the operating system, such as the kernel.

The interesting thing about containers is the application itself. The application might have various dependencies it requires to run. These dependencies exist only within the container itself. This means that if something bad that happens to Application A and the binaries it depends on, it has no impact on Application B and the binaries on which it depends. For example, in most environments, if you delete the registry from Application A, the consequences are disastrous for both Application A and Application B. However, with containers, Application A and Application B are each self-contained, and the change to the registry for Application A does not affect Application B.

Because all binaries and dependencies are hosted within the container, the application running in the container is completely portable. Essentially, this means that you can deploy a container to any host running the container manager software, and it will start and run without any modification.

Containers are built on layers. The first layer is the *base layer*. This is the operating system image on which all other layers will be built. This image is stored in an image repository so that you can reference it when necessary. The next layer (and sometimes the final layer) is the *application framework layer* that can be shared between all of your applications. For example, if your base layer is Windows Server Core, your application framework layer could be .NET Framework and Internet Information Services (IIS). The second layer can also be stored as an image, which, when called, also describes its dependency on the base layer of Windows Server Core. Finally, the *application layer* is where the application itself is stored, with references to the application framework layer and, in turn, to the base layer.

The base layer and the application layer can be referenced at any time by any other application container you create. Each layer is considered read-only except the top layer of the “image” you are deploying. For example, if you deploy a container that depends only on the Windows Server Core image, this Windows Server Core layer is the top layer of the container and a sandbox is put in place to store all the writes and changes made during runtime. You can then store the changes made as another image for later reuse. The same applies if you deploy the application framework layer image; this layer would have its own sandbox, and if you deploy your application to it, you can then save the sandbox as a reusable image.

Basically, when you deploy a container to a host, the host determines whether it has the base layer. If not, it pulls the base layer from an image repository. Next, it repeats the process for the application framework layer and then creates the application container you were originally trying to deploy. If you then want to create another container with the same dependencies, you simply issue a command to

create the new application container, and it is provisioned almost immediately because all of the dependencies are already in place. If you have an application container that depends on a different application framework layer as well as on the original Windows Server Core base layer, you can simply pull the different application framework layer from an image store and start the new application container.

Why use containers?

Containers provide some distinct advantages over the traditional model of deploying an application into a VM or onto a physical host.

The first advantage has to do with development. A general pain for developers when they are building applications revolves around moving the application from a development environment, to test, and then to production. Developers must spend a lot of time and effort checking the application's dependencies as it moves through the environments. However, when an application is deployed to a container, you can move the container between environments because it is isolated and all binaries are self-contained within the container itself.

Another reason for using containers is to achieve higher scale versus deploying an application to a VM. To achieve the different environments of development, test, and production in a VM model, you need at least three VMs; in a container model you need only one. That single VM, running a container manager, can run three containers simulating development, test, and production environments. With containers, you need fewer VMs to run your environments and you can achieve significantly higher scale in your cloud environments.

Containers also allow for rapid deployment and operation of applications. Unlike VMs, containers don't have an underlying operating system, as such. Think in terms of deployment. If you want to create a new application or scale the existing application to support more load, you just load a new container; the operating system is already in place. This means that the time spent waiting for a container to deploy or scale up is significantly shorter than with a VM because you are never waiting for the operating system to start up.

Windows Server containers versus Hyper-V containers

Two types of containers will be available in Windows Server 2016 Technical Preview:

- Windows Server containers
- Hyper-V containers

You can consider Windows Server containers to be the equivalent to Linux containers. Windows Server container types isolate applications on the same container host. Each container has its own view of the host system, including the kernel, processes, file systems, the registry, and other components. In the case of Windows Server containers, they work between the user-mode level and the kernel-mode level.

Hyper-V containers are based on a container technology that is rooted in hardware-assisted virtualization. With hardware-assisted virtualization, Hyper-V containers' applications are provided a highly isolated environment in which to operate, where the host operating system cannot be affected in any way by any running container.

Figure 6-7 shows what the layout might look like in relation to the two container technologies that are available in Windows Server 2016 Technical Preview.




Figure 6-7: Windows Server 2016 containers and Hyper-V containers on the same physical box

As shown in the illustration, one physical host can offer a mix of Hyper-V containers, VMs, and Windows Server containers. With two potential container options, the question is when to use either Windows Server containers or Hyper-V containers for developing and deploying the applications. The deciding criteria come when you identify scale and certified, hardware-assisted isolation requirements for the application or customer using the containers.

For example, if you require greater scale, Windows Server containers are the best option because you can achieve this goal. However, if you require the isolation benefits of hardware-assisted virtualization, and scale is not as important, Hyper-V containers are the best option to use.

When looking at either type of containers, it is important to understand the development lifecycle. From a developer's perspective, the tools they are familiar with (such as Visual Studio) will make it possible for them to write and deploy applications directly into containers. The tools will also give them the ability to describe what core functionality is required in the underlying operating system, what libraries can be shared between applications, or what libraries are to be dedicated for a container. Figure 6-8 highlights the dependencies and runtime mode of the containers.




Figure 6-8: Containers and dependencies

No matter which technology you choose, the application you deploy into a container is compatible between both technologies. This essentially means that a developer can easily build the application in a container hosted on Windows Server containers and move it to a Hyper-V container with no changes required. This gives great flexibility, especially if the requirements of scale or isolation change after the initial planning of the system.

So, how do you use or get started with containers in Windows Server 2016 Technical Preview? First, you need to understand that the general process to work with containers:

- Turn on the Windows “containers” feature
- Create a VM switch
- [Optional] Configure Network Address Translation (NAT) if required
- Install a container OS image
- [Optional] Deploy Docker
- [Optional] Turn on Hyper-V
- [Optional] Turn on nested virtualization
- [Optional] Configure VP’s for nested VM
- [Optional] Turn off dynamic memory for nested VM
- [Optional] Turn on MAC spoofing for nested VM

As you can see, there are a variety of optional components, and it comes down to how you want to use containers and what technologies you want to use. For example, if you don’t need the isolation techniques described earlier in this section, you do not need to deploy Hyper-V.

In the following example, you are going to install Windows Server 2016 Technical Preview containers. The first thing you need to do is install the Windows Feature to turn on Container Support.

Using Windows PowerShell, you can use the following command to install the containers feature:

```
Install-WindowsFeature containers
```

After the components have been installed, restart the machine to ensure that they are fully turned on.

You can verify the installation after the restart by using the Get-ContainerHost cmdlet:

```
get-containerhost
```

```
Name      ContainerImageRepositoryLocation
-----
CLOUD01 C:\ProgramData\Microsoft\Windows\Hyper-V\Container Image Store
```

Next, you need to give the container images a method by which to connect to the outside world. Using the New-VMSwitch cmdlet, you can create a VM switch that you can use to give container images a view to the outside world. However, you do have a choice to make before we do this: You can configure a container VM switch as NAT or external. The external switch type makes it possible for you to configure your container image to be assigned an IP address directly from your corporate network or receive an IP address from DHCP. Alternatively, if you want containers to have their own virtual network that doesn’t overlap with your corporate network, you can configure NAT to hide the network. Then, you can open endpoints to the container images and only expose the services you require.

The following example creates an external switch. Note that you need to know the name of the adapter to which you want to bind the external VM switch.

```
New-vmswitch -Name ContainerSW -NetAdapterName Ethernet
```

Alternatively, for a NAT-configured switch, you can use the following example. Note that in this case you don’t need an external adapter to create the switch.

```
New-VMSwitch -Name "NAT" -SwitchType NAT -NATSubnetAddress 192.168.0.0/20
```

If you create a VM switch, you need to create an object which will allow the translation to happen. This is called the *NAT Object*, and you can use the following example to create it:

```
New-NetNat -Name 'NAT' -InternalIPInterfaceAddressPrefix '192.168.0.0/20'
```

Next, you need to get some container images that you can use as the basis for your images moving forward. Base container images are available from a source repository and can be installed to your host. First you need to install a provider which will make it possible for you to see the information in the repository. Using Windows PowerShell, you can use the `Install-PackageProvider` cmdlet to install the container provider, as follows:

```
Install-PackageProvider ContainerProvider -Force
```

Use the `Find-ContainerImage` cmdlet to browse the public repository of container images available. The `Find-ContainerImage` cmdlet uses the Windows PowerShell OneGet package manager in the background to retrieve the listings. To avoid confusion, ensure that you select only the container image for the host type to which you are deploying it. For example, don't download and install a Nano Server image on a Windows Server Hyper-V Host, but you can have a nested Nano Server and download a Nano Server image to that nest container host.

The following example shows the cmdlet `Find-ContainerImage` and its output:

```
Find-ContainerImage
```

Name	Version	Description
NanoServer	10.0.10586.0	Container OS Image of Windows Se...
WindowsServerCore	10.0.10586.0	Container OS Image of Windows Se...

Now, you can select and install the image that you want by using the `Install-ContainerImage` cmdlet, as follows:

```
Install-ContainerImage -Name NanoServer -Version 10.0.10586.0
```

After the image installation is finished, use the `Get-ContainerImage` cmdlet to verify the installation.

```
get-containerimage
```

Name	Publisher	Version	IsOSImage
NanoServer	CN=Microsoft	10.0.10586.0	True

The next step is to deploy a container from this container image. Using the `New-Container` cmdlet, you can build your first container, as shown in the following example:

```
$container = get-containerimage -name "NanoServer"  
New-Container -ContainerImage $container -Name W2016C1 -ContainerComputerName W2016C1
```

Name	State	Uptime	ParentImageName
W2016C1	Off	00:00:00	NanoServer

When the container is first built, it is off, which is good because you also have no network bound yet. So, using the `Add-ContainerNetworkAdapter`, you can create a network card in the container.

```
Add-ContainerNetworkAdapter -ContainerName W2016C1
```

Then, you can use the `Connect-ContainerNetworkAdapter` cmdlet to attach the NIC to the switch.

```
Connect-ContainerNetworkAdapter -ContainerName W2016C1 -SwitchName NAT
```

First, store the new container into a variable and then start the container by using the `Start-Container` cmdlet.

```
$container = get-container -name W2016C1  
Start-Container $container
```

To stop the container, you can use the Stop-Container cmdlet.

The container is now up and running, and you obviously would like to manage what is happening within it. You can use the Enter-PSSession cmdlet with a new parameter ContainerName to start a Remote PSSession to the container.

```
Enter-PSSession -ContainerName W2016C1
```

The session is started then with the container you have running; for example, you can run an IPConfig and validate that you are indeed in the container and running on the right IP address space.

Note Container cmdlets and support are continuously increasing as we approach general availability of Windows Server 2016. For updated information and samples about working with containers, go to https://msdn.microsoft.com/virtualization/windowscontainers/containers_welcome.

What about Docker?

Containers are not a new technology. In general, they have existed in the Linux world for quite some time. Docker is an open-source engine that has helped containers become more prevalent.

Currently, Docker's open-source runtime builds, ships, and runs containers on Linux operating systems. Because it is open source, an extensive ecosystem of developers, and now "dockerized applications," have grown up around it. Docker provides a user-friendly experience to manage the lifecycle of its containers, facilitating easy adoption.

In 2015, Microsoft Azure announced support for the Docker engine on Linux VMs in Azure. This was exciting news, but the question remains: What about Docker on Windows natively? With the introduction of Windows Server containers and Hyper-V containers, Docker becomes even more useful because you can use it to manage Docker containers on Windows as well as the traditional Linux environment. Also, we now have access to all of the images that are available through Docker, so we can download and deploy!

The Docker runtime engine will work as an abstraction on top of Windows Server containers and Hyper-V containers. Docker provides all the necessary tooling to develop and operate its engine on top of Windows containers, be it Hyper-V containers or Windows Server containers. This will afford the same flexibility of developing an application in one container and being able to truly run it anywhere.

Figure 6-9 shows the placement of the Docker engine in relation to Windows Server containers or Hyper-V containers and compares it to a Docker engine running on Linux.




Figure 6-9: The Docker engine on Windows and Linux

The Docker engine runs at the same level in either a Windows Server container or Linux container environment, and it can run with Windows Server or Linux above the Docker engine. The Docker client will connect to any Docker engine and provide a consistent management experience for the end user.

For Docker to manage your container environment, you first need the Containers feature installed. Otherwise, you will not be able to manage the Windows Server or Hyper-V containers environment with Docker.

Docker is a third-party product and doesn't come installed or available as a feature in Windows. We must download the installer/binaries and set up our host with the Docker Engine. From Windows PowerShell, you can download the Docker Engine as follows by using wget:

```
wget https://aka.ms/tp4/docker -OutFile $env:SystemRoot\system32\docker.exe
```

The Docker Engine needs to be placed in the \$env:SystemRoot\System32 directory on Windows Server or on Nano Server, or both, depending on which one you are using as your Windows container host.

After it has downloaded, it is important to note that Docker is a self-contained executable; it does not have an installer. As of this writing, it does not install as a Windows service. Therefore, you must start the Docker Engine every time the host is restarted. You can use a batch file or Windows PowerShell script on startup or a third-party tool, the Non-Sucking Service Manager (NSSM), to install a service to do this for you.

Whichever way you choose, Docker requires some syntax to start its engine. The syntax will vary based on whether you choose to run Docker in secure mode or unsecure mode. The secure mode requires certificate authentication to manage the Docker environment, limiting the amount of hosts/management clients that can connect.

For unsecure mode, the syntax is as follows:

```
Docker daemon -D -b '<switchname>'
For secure mode the syntax is as follows
Docker daemon -D -b '<switchname>' -H <ipaddress>:2376 -tlsverify -tlscacert=<certpath>\<rootcertname.pem>
-tlscert=<certpath>\<servercertname>.pem
-tlskey=<certpath>\<servercertkey>.pem
```

- **-D** stands for Debug mode
- **-b** stands for bridge (i.e., connect to your virtual switch)
- **-H** stands for Host

When the engine is up and running, you can open an additional command prompt window and use the Docker.exe file to interact with the overall session. For example, if you want to list the images installed on the box, you can call Docker images, as follows:

docker images				
REPOSITORY	TAG	IMAGE ID	CREATED	VIRTUAL SIZE
windowsservercore	10.0.10586.0	6801d964fda5	4 months ago	0 B
nanoserver	10.0.10586.0	8572198a60f1	4 months ago	0 B

As you can see, these are the images you downloaded and installed in your Windows Server container environment.

To start a container using Docker, you use the Docker run command, as follows:

```
Docker run - -name TestContainer1 -it windowsservercore cmd
```

This will invoke a container which will start a command session when it's deployed, and you can interact with the image.

The principal from there is you can work with that image, customize it to your needs, and save it for later use or redeployment.

More info This book does not present an in-depth view of Docker. For more information, go to https://msdn.microsoft.com/virtualization/windowscontainers/quick_start/manage_docker.

Systems management

This chapter explores some of the new elements for Windows Server 2016 Technical Preview related to systems management. The core areas we will discuss are Windows PowerShell V5 and then we will dive into detailing what's new in System Center 2016 and how you can take advantage of Microsoft Operations Management Suite to have a complete hybrid management experience.

Windows PowerShell improvements

By Ritesh Modi

Windows PowerShell Desired State Configuration (DSC) is a hot topic nowadays. DSC is a new management platform with which administrators can use Windows PowerShell for deploying and managing configuration data for software services and also for managing the environment in which these services run. Windows Server 2016 Technical Preview introduces several improvements to DSC, and in this chapter, Ritesh Modi examines two of these enhancements: the new Local Configuration Manager v2 and a new partial configuration feature. Ritesh also examines the new PowershellGet module and the NuGet package manager, two of the many new capabilities that you will find in Windows PowerShell v5.

DSC Local Configuration Manager

One of the most important features of DSC is Local Configuration Manager (LCM) for both DSC push and pull architecture. It acts as the DSC client engine responsible for accepting (push mode) or retrieving (pull mode) configurations (MOF files), running MOF files, monitoring, comparing the MOF configuration with current server configuration, reporting the drifts, and applying (or reapplying) the configurations. Needless to say, LCM is the heart and brain of DSC, and it should be installed on all the servers that are managed by DSC within a network.

Windows Server 2016 Technical Preview comes preinstalled with Windows Management Framework 5.0 and DSC. DSC was introduced with Windows Management Framework 4.0 and was part of Windows Server 2012 R2 and Windows 8.1. As you might expect, Windows Management Framework 5.0 introduces many new features and changes to DSC and LCM.

LCM has configurable settings known as Meta Configuration. The behavior and actions of LCM can be influenced and controlled by modifying Meta Configuration properties.

DSC v2 in Windows Server 2016 Technical Preview builds on the previous version, deprecating a few properties and offering newer cmdlets for managing configurations, a newer LCM with additional functionality, newer Meta Configuration attributes and properties, and new features such as partial configurations and cross-computer synchronization.

LCM is implemented as a Common Information Model (CIM) class `MSFT_DSCLocalConfigurationManager` within the `root\Microsoft\Windows\DesiredStateConfiguration` namespace.

In this section, we look into the newer functionality and workings of LCM v2 on Windows Server 2016 Technical Preview.

DSC provides two cmdlets for accessing LCM and viewing and updating the Meta Configuration properties: `Get-DSCLocalConfigurationManager` and `Set-DSCLocalConfigurationManager`. Running `Get-DSCLocalConfigurationManager` on the Windows Server 2016 Technical Preview Windows PowerShell console lists all LCM Meta Configuration properties along with their current values. The default values are shown in the following output:

```
PS C:\> Get-DscLocalConfigurationManager
```

```
AllowModuleOverWrite      : False
CertificateID             : 
ConfigurationDownloadManagers : {}
ConfigurationID           : 
ConfigurationMode         : ApplyAndMonitor
ConfigurationModeFrequencyMins : 15
Credential                : 
DebugMode                 : False
DownloadManagerCustomData : 
DownloadManagerName       : 
LCMCompatibleVersions     : {1.0, 2.0}
LCMState                  : Ready
LCMVersion                 : 2.0
MaxPendingConfigRetryCount : 
StatusRetentionTimeInDays : 7
PartialConfigurations     : {}
RebootNodeIfNeeded        : False
RefreshFrequencyMins       : 30
RefreshMode               : PUSH
ReportManagers             : {}
ResourceModuleManagers    : {}
PSComputerName
```

LCM v2 in Windows Server 2016 Technical Preview includes all of the properties from the first version for backward compatibility. Some of these properties are needed in the newer version, whereas others are deprecated and cannot be used for configuring LCM v2. The following is a summary of the LCM properties in Windows Server 2016 Technical Preview:

- **ConfigurationModeFrequencyMins** Used in pull mode. This represents how frequently in minutes LCM should request configuration from the pull server. The default is 15 minutes.
- **RebootNodeIfNeeded** Used for signifying whether the server needs to reboot after any action performed by DSC resources.
- **ConfigurationID** Used in pull mode. The value is a globally unique identifier (GUID) representing a configuration document on the pull server.

- **ConfigurationMode** Valid values are ApplyOnly, ApplyandMonitor, and ApplyandAutoCorrect. The default is ApplyandMonitor.
- **RefreshFrequencyMins** Represents how frequently in minutes LCM should check and/or apply configuration. The default is 30 minutes.
- **AllowModuleOverWrite** Used in pull mode. Determines whether dynamically downloaded DSC resources should overwrite existing resources with the same name.
- **CertificateID** Deprecated in LCM v2. This information is now captured using other properties.
- **Credential** Deprecated in LCM v2. This information is now captured using other properties.
- **DownloadManagerName** Deprecated in LCM v2. This information is now captured using other properties.
- **DownloadManagerCustomData** Deprecated in LCM v2. This information is now captured using other properties.

Compared to the previous version, LCM v2 includes a number of new properties. Some of the important new properties in LCM v2 include the following:

- **DebugMode** When the value for this property is True, it causes the engine to reload the PowerShell DSC resource.
- **ConfigurationDownloadManagers** This property replaces the DownloadManagerName and DownloadManagerCustomData properties. It indicates whether a connection is secure and what certificate to use when connecting to the pull server. Valid values are ConfigurationRepositoryWeb and ConfigurationRepositoryShare.
- **ResourceModuleManagers** Represents the location for downloading DSC resources. Valid values are MSFT_WebResourceManager and MSFT_FileResourceManager.
- **ReportManagers** Represents the details of the compliance server. The valid value is MSFT_WebReportManager.
- **PartialConfigurations** Represents an array of configuration fragments that together form complete configuration for the node.
- **LCMUpdatename** Used internally by DSC.
- **StatusRetentioninTimeinDays** By default, every 15 minutes the configuration is applied by LCM and its status is stored on the file system. This setting determines how long the status files are retained. The default value is seven days. This means status files are retained for seven days, after which they are deleted.

With the new LCM properties, it is possible to have multiple configuration fragments instead of a single configuration. The properties are more organized, each with well-defined usage. You can query the current state of LCM, Turn on and turn off caching, and separate URL endpoints for configurations and resources.

DSC v1 used a special resource, LocalConfigurationManager, to set the LCM Meta Configuration properties. This resource is deprecated in LCM v2. It can still be used to configure LCM v2; however, it cannot configure new Meta Configuration properties. It is recommended that you use the new Settings resource to set LCM properties, instead.

You set MetaConfiguration properties by performing a few steps in sequence. As mentioned, in LCM v2, you should use a new special resource named Settings to configure LCM Meta Configuration properties. You should place this new resource in a configuration script and run it. Running the

resource generates a MOF file, which is sent to the LCM of the destination server. The LCM of the destination server applies and changes Meta Configuration property values. Note that LCM configuration is not allowed in a regular configuration comprising general DSC resources. Along with the Settings resource, a few more LCM-specific resources are included in LCM v2. These LCM resources provide a better authoring experience and eventually change the properties available in the Settings resource. They are summarized as follows:

- **Settings** This is the primary LCM Meta Configuration resource.
- **ConfigurationRepositoryWeb** This resource represents the Internet Information Services (IIS) Open Data Protocol (OData) endpoint for pull servers. This resource changes the ConfigurationDownloadManagers property of the Settings resource. It has the following properties:
 - Name
 - ServerUrl
 - AllowUnsecureConnection
 - CertificateID
- **ConfigurationRepositoryShare** This resource represents the Server Message Block (SMB) share endpoint for pull servers. This resource changes the ConfigurationDownloadManagers property of the Settings resource. It has the following properties:
 - Name
 - SourcePath
 - Credential
- **MSFT_WebResourceManager** This resource represents the IIS OData endpoint for downloading DSC resources by using IIS. This resource changes the ResourceModuleManagers property of the Settings resource. It has the following properties:
 - Name
 - ServerUrl
 - CertificateID
 - Priority
- **MSFT_FileResourceManager** This resource represents the IIS OData endpoint for downloading DSC resources by using SMB shares. This resource changes the ResourceModuleManagers property of the Settings resource. It has the following properties:
 - Name
 - ServerUNC
 - Credential
 - Priority

- **MSFT_WebReportManager** This resource represents the IIS OData endpoint for providing reporting data related to nodes, their current configurations, and drifts. This resource changes the ResourceModuleManagers property of the Settings resource. It has the following properties:
 - Name
 - ServerUrl
 - CertificateID
 - AllowUnsecureConnection
 - CustomData
- **PartialConfiguration** This resource represents the name of the configuration that should be pulled from the pull server. It is possible to have multiple PartialConfiguration resources in a configuration. This resource changes the PartialConfigurations property of the Settings resource. It has the following properties:
 - Description
 - ExclusiveResources
 - ConfigurationSource

The following is a typical implementation of the Meta Configuration property in LCM v2:

```
[DSCLocalConfigurationManager()]
Configuration ChangeLCMProperties
{
    Node DemoServerWin
    {
        Settings
        {
            AllowModuleOverwrite = $false
            RebootNodeIfNeeded = $true
            RefreshMode = "Pull"
            ConfigurationMode = "ApplyAndAutoCorrect"
            ConfigurationID = "fcd03a8d-5a64-4982-92b3-5c89680add39"
        }

        ConfigurationRepositoryWeb PullServer1
        {
            Name = "PullServer1"
            ServerURL = "http://demoServerwin10:8090/PSDSCPullServer.svc/"
            AllowUnsecureConnection = $true
        }

        ConfigurationRepositoryWeb PullServer2
        {
            Name = "PullServer2"
            ServerURL = "http://demoServerwin10:8080/PSDSCPullServer.svc/"
            AllowUnsecureConnection = $true
        }

        MSFT_WebReportManager ComplianceServer
        {
            Name = "ComplianceServer"
            ServerURL = "http://demoServerwin10:8080/PSDSCComplianceServer.svc/"
            AllowUnsecureConnection = $true
        }

        PartialConfiguration IISInstall
        {
            Description = 'Configuration for IIS Web Server'
            ConfigurationSource = '[ConfigurationRepositoryWeb]PullServer1'
        }

        PartialConfiguration IndexFile
        {
            Description = 'Configuration for Index File'
            ConfigurationSource = '[ConfigurationRepositoryWeb]PullServer2'
            DependsOn = '[PartialConfiguration]IISInstall'
        }
    }
}
```

```

    }
}
}

ChangeLCMProperties -OutputPath "C:\DSC"

Set-DscLocalConfigurationManager -Path "C:\DSC"

```

The preceding script is similar to general DSC configuration with the configuration named `ChangeLCMProperties` but including the defined `DSCLocalConfigurationManager` attribute. This attribute mandates that all resources within the configuration should be related to LCM only and should be present on configurations related to LCM. An error results if other general resources are used in the configuration. The script contains one node section for the `DemoServerWin` server.

The `Settings` resource is the main resource for setting the LCM properties. In this example, we are specifying some of its properties and assigning values to them. For example, the refresh mode is set to `Pull` so that the machine should restart (when required by a resource); the configuration mode has been changed to `ApplyAndAutoCorrect`, and `ConfigurationID` has been provided with a valid GUID. The configuration as represented by the GUID would be pulled from the pull server.

There are two pull servers in this configuration denoted by `PullServer1` and `PullServer2`. The `ServerURL` property shows that they are on the same server with different port numbers. Also, `AllowUnsecureConnection` makes it possible to use HTTP instead of HTTPS protocol. Compliance server information is also provided by using `MSFT_WebReportManager`. There are two partial configurations set to be downloaded by LCM and applied as a single configuration on its server. Partial Configuration `IISInstall` is responsible for downloading a configuration named `IISInstall` from Pull Server1. Partial Configuration `IndexFile` is responsible for downloading a configuration named `IndexFile` from Pull Server2. Moreover, running the partial configuration `IndexFile` depends on the completion of the configuration `IISInstall` as represented by the `DependsOn` property. Only after the `IISInstall` configuration is applied can the `IndexFile` configuration run.

After the configuration is defined, it runs to generate the MOF file (`DemoServerWin.Meta.mof`) at `C:\DSC`. The folder location has been explicitly provided by using the `OutputPath` parameter. After the MOF file is generated, the `Set-DSCLocalConfigurationManager` cmdlet is used to push and apply the MOF file to the `DemoServerWin` server.

When the configuration is applied, LCM on the server `DemoServerWin` is configured to pull partial configurations from multiple pull servers and apply them periodically.

Next, the new configuration can be read again by using `Get-DSCLocalConfigurationManager`, as shown in the following example:

```

PS C:\Users\me> Get-DscLocalConfigurationManager

AllowModuleOverWrite           : False
CertificateID                   :
ConfigurationDownloadManagers  : {[ConfigurationRepositoryWeb]PullServer1,
                                   [ConfigurationRepositoryWeb]PullServer2}
ConfigurationID                 : fcd03a8d-5a64-4982-92b3-5c89680add39
ConfigurationMode               : ApplyAndAutoCorrect
ConfigurationModeFrequencyMins : 15
Credential                     :
DebugMode                      : False
DownloadManagerCustomData      :
DownloadManagerName            :
LCMCompatibleVersions          : {1.0, 2.0}
LCMState                       : Ready
LCMVersion                    : 2.0
MaxPendingConfigRetryCount     :
StatusRetentionTimeInDays      : 7
PartialConfigurations          : {[PartialConfiguration]IISInstall,
                                   [PartialConfiguration]IndexFile}
RebootNodeIfNeeded             : True
RefreshFrequencyMins           : 30
RefreshMode                    : PUSH

```

```
ReportManagers           : [MSFT_WebReportManager]ComplianceServer
ResourceModuleManagers   : {}
PSComputerName           :
```

In the preceding code block, the ConfigurationDownloadManagers property is filled with two values representing two pull servers: PartialConfigurations has two values represented by the IISInstall and IndexFile configurations, and ReportManagers has a value of ComplianceServer.

New methods in LCM

There are three new methods in LCM v2: GetConfigurationStatus, GetConfigurationResultOutput, and SendConfigurationApplyAsync. Let's examine these briefly.

GetConfigurationStatus

The GetConfigurationStatus method retrieves the current status of configuration for a server. The new DSC cmdlet Get-DSCConfigurationStatus invokes the CIM method. The following code shows an example of Get-DSCConfigurationStatus:

```
PS C:\> $cimsession = New-CimSession -ComputerName DemoServerWin10
PS C:\> Get-DscConfigurationStatus -CimSession $cimsession
```

Status	StartDate	Type	Mode	RebootRequested
	NumberOfConfigurationResources	PSComputerName		
Success	2014/11/12 16:23:03	Consistency	PUSH	False
1		DemoServerWin10		

GetConfigurationResultOutput

The GetConfigurationResultOutput method provides verbose information about the current configuration and the configuration drifts. There is no DSC cmdlet that invokes this CIM method. Instead, you can invoke it can by using the CIM cmdlet Invoke-CIMMethod, as shown here:

```
PS C:\>
$ConsistencyCheck = (Invoke-CimMethod -ClassName "MSFT_DSCLocalConfigurationManager" `
    -Namespace "root\Microsoft\Windows\DesiredStateConfiguration" `
    -MethodName getConfigurationResultOutput)

for($i=0; $i -le 100; $i++)
{
    $ConsistencyCheck[$i].ItemValue.Message
}
[DEMOSERVERWIN10]: [] Starting consistency engine.
[DEMOSERVERWIN10]: LCM: [ Start Resource ] [[WindowsFeature]XPS]
[DEMOSERVERWIN10]: LCM: [ Start Test ] [[WindowsFeature]XPS]
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] Begin running Test functionality on the
XPS-Viewer feature.
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] Querying for feature XPS-Viewer using
Server Manager cmdlet Get-WindowsFeature.
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] The operation 'Get-WindowsFeature'
started: XPS-Viewer
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] GetServerComponentsAsync provider method
started: XPS-Viewer
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] Call to GetServerComponentsAsync provider
method succeeded.
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] The operation 'Get-WindowsFeature'
succeeded: XPS-Viewer
[DEMOSERVERWIN10]: [[WindowsFeature]XPS] End running Test functionality on the
XPS-Viewer feature.
[DEMOSERVERWIN10]: LCM: [ End Test ] [[WindowsFeature]XPS] in 0.4667 seconds.
[DEMOSERVERWIN10]: LCM: [ End Resource ] [[WindowsFeature]XPS]
[DEMOSERVERWIN10]: [] Consistency check completed.
```

SendConfigurationApplyAsync

The SendConfigurationApplyAsync method applies the configuration to a target server asynchronously. This means LCM invokes this method and does not wait for its completion. Again, there is no DSC cmdlet to invoke this method; however, you can invoke it through a CIM cmdlet, as shown in the following example:

```
PS C:\> Configuration PushDemo
{
    Node DemoServerWin10
    {
        WindowsFeature XPS
        {
            Name = "XPS-Viewer"
            Ensure = "Absent"
        }
    }
}

PushDemo -OutputPath "C:\DSC"
$mofString = get-content "C:\dsc\DemoServerWin10.mof"
$mofbytes = [System.Text.Encoding]::ASCII.GetBytes($mofString)
$AsyncApply = Invoke-CimMethod -ClassName "MSFT_DSCLocalConfigurationManager" `
    -Namespace "root\Microsoft\Windows\DesiredStateConfiguration" `
    -MethodName SendConfigurationApplyAsync `
    -Arguments @{ConfigurationData=$mofbytes;Force=$true}

$AsyncApply

Directory: C:\DSC
Mode                LastWriteTime         Length Name
----                -
-a-----          11/12/2014    2:28 PM           1226 DemoServerWin10.mof
PSComputerName :
```

DSC partial configurations

One of the most awaited and interesting features of DSC v2 is partial configuration. Until DSC v2, it was difficult to split a configuration into multiple configuration files authored for a server. Partial configuration makes it possible for you to split a configuration into multiple smaller configuration fragments across multiple files. Partial configurations are implemented exactly the same as any general DSC configuration. It is the responsibility of LCM on a destination server to combine all the configuration fragments into a single configuration and apply it.

Partial configurations are complete in and of themselves and can be applied independently as a complete configuration to any server. It is the way they are deployed on a pull server and the way the LCM Meta Configuration is configured on the target server that makes it possible for partial configurations to be applied to a server.

In Windows Server 2016 Technical Preview, partial configurations only work with DSC pull mode. This means that you should configure the LCM of servers in a network to pull configurations from a pull server (IIS or SMB share) and be able to identify the configurations distinctly on these pull servers.

The benefits of partial configurations include the following:

- Multiple authors can author configurations independently and simultaneously for servers in a network.
- You can apply incremental configurations to servers without modifying any existing configurations.
- Modular authoring of configurations is available.
- There are no longer dependencies on using only a single MOF file. This was the case in DSC v1, for which only one MOF file was allowed and applied to a server at a given point of time. Newer configuration (MOF) would replace the current configuration in DSC v1.

To make partial configuration work in Windows Server 2016 Technical Preview, complete the following steps:

1. Create the pull server.
2. Configure the LCM Meta Configuration of servers on the network.
3. Author the configurations.
4. Deploy the configurations on the pull server.

We will look into the details of each of these steps, except for the creation of the pull server because that process is the same as for DSC v1.

Setting up the LCM Meta Configuration

To prepare a server's LCM Meta Configuration, you must set the following:

- RefreshMode with the value of Pull.
- ConfigurationID with the value of a valid GUID, representing a configuration of a pull server. (This GUID is required on a pull server for naming the configuration files.)
- ConfigurationMode with the value of ApplyandAutoCorrect to keep the server in the expected state.
- Multiple WebConfigurationRepository resource instances, each representing a pull server.
- Multiple PartialConfiguration resource instances, each representing a configuration on a pull server.

To demonstrate partial configuration, the following example features an environment with two pull servers (DemoServerWin10 and ServerWin10). There are also two configurations, each deployed to one of the pull servers. The LCM of a destination machine is configured with these two pull servers and configurations. All of the steps for partial configuration are done on these servers. The LCM configuration applied to the server named DemoServerWin is shown in the following:

```
[DSCLocalConfigurationManager()]
Configuration ChangeLCMProperties
{
    Node DemoServerWin
    {
        Settings
        {
            RebootNodeIfNeeded = $true
            RefreshMode = "Pull"
            ConfigurationMode = "ApplyAndAutoCorrect"
            ConfigurationID = "fcd03a8d-5a64-4982-92b3-5c89680add39"
        }

        ConfigurationRepositoryWeb PullServer1
        {
            Name = "PullServer1"
            ServerURL = "http://serverwin10:9000/PSDSCPullServer.svc/"
            AllowUnsecureConnection = $true
        }

        ConfigurationRepositoryWeb PullServer2
        {
            Name = "PullServer2"
            ServerURL = "http://demoServerwin10:8080/PSDSCPullServer.svc/"
            AllowUnsecureConnection = $true
        }

        PartialConfiguration IISInstall
        {
            Description = 'Configuration for IIS Web Server'
            ConfigurationSource = '[ConfigurationRepositoryWeb]PullServer1'
        }
    }
}
```



```

    }

    PartialConfiguration IndexFile
    {
        Description          = 'Configuration for Index File'
        ConfigurationSource = '[ConfigurationRepositoryWeb]PullServer2'
        DependsOn            = '[PartialConfiguration]IISInstall'
    }
}

ChangeLCMProperties -OutputPath "C:\DSC"
Set-DscLocalConfigurationManager -Path "C:\DSC"

```

Assuming that the pull servers are already available on the mentioned servers, you must change the LCM Meta Configuration settings on all servers that will participate and pull configurations from them. The preceding configuration sample has a node section named `DemoServerWin`. It signifies that the configuration modifies the LCM configuration of `DemoServerWin`. Attribute `DSCLocalConfigurationManager` mandates that only resources applicable for LCM Meta Configuration can be used in this configuration. General resources cannot be used in such configurations. Using this attribute is the way to indicate to DSC that this configuration relates to LCM configuration.

The `Settings` resource is configured with the `RefreshMode` property set to `Pull`, the `ConfigurationMode` property set to `ApplyandAutoCorrect`, the `ConfigurationID` property set to `fcd03a8d-5a64-4982-92b3-5c89680add39`, and the `RebootNodeIfNeeded` property set to `True`. LCM downloads configuration files from the pull server whose name has the same GUID as that assigned to the `ConfigurationID`. In Listing 5-7, LCM would download configuration files with `fcd03a8d-5a64-4982-92b3-5c89680add39` in their names from all pull servers.

You also need to provide the pull server details to LCM. In LCM v2, you can do so by using the `WebConfigurationRepository` resource. There can be multiple pull servers (`WebConfigurationRepository` resources) defined in a configuration. In the previous sample, two pull servers are defined: `PullServer1` with server URL `http://serverwin10:9000/PSDSCPullServer.svc/` and `AllowUnsecureConnection` set to `True`, and `PullServer2` with server URL `http://Demoserverwin10:8090/PSDSCPullServer.svc/` and `AllowUnsecureConnection` set to `True`. `AllowUnsecureConnection` allows LCM to request configuration on HTTP protocol instead of HTTPS protocol.

The `PartialConfiguration` resource defines configuration fragments. Two partial configurations, `IISInstall` and `IndexFile`, are defined. `IISInstall` configuration is available on `PullServer1`, whereas `IndexFile` configuration is available on `PullServer2`. Important to note are the names of the partial configurations because they should exactly match the names of the configurations on the pull server. The next section will show that `IISInstall` configuration is authored and available on `PullServer1` and `IndexFile` configuration is available on `PullServer2`. The `ConfigurationSource` property attaches the pull server to the partial configuration.

Also, note that the pull server URL, `ConfigurationID`, and `Configuration Name` combined provide LCM with complete information to uniquely identify configuration on the pull server. LCM cannot pull partial configurations if any of these three pieces of information is missing.

The configuration previously shown generates the `DemoServerWin.meta.mof` file at the `C:\DSC` folder location. You use the `Set-DSCLocalConfigurationManager` cmdlet to push and apply the MOF file to `DemoServerWin`.

Authoring the configurations

Next, we examine authoring the configurations that will participate in partial configuration. In this section, two configurations, `IISInstall` and `IndexFile`, are authored on separate servers, `DemoServerWin10` and `ServerWin10`, respectively.

IISInstall configuration

This is a simple configuration responsible for installing IIS (Web-Server) on a server using the WindowsFeature resource. Running the configurations in this section and the next one generates MOF files. The configuration script in our sample that follows runs on server ServerWin10. The name of the MOF file is same as the node name defined in the configuration script. Configurations participating in partial configurations have a special naming requirement. They should use the <<ConfigurationName>>.<<ConfigurationID>>.MOF format. The configuration shown in the sample that follows uses ConfigurationData (data structure for passing values to configuration script) to define the name of the node. \$AllNodes.NodeName retrieves all the node names from the configuration data, which in this case is only one because there is just one NodeName. IISInstall.fcd03a8d-5a64-4982-92b3-5c89680add39.MOF is generated from the following script:

```
$ConfigInfoIIS = @{
    AllNodes = @(
        @{
            NodeName = "IISInstall.fcd03a8d-5a64-4982-92b3-5c89680add39"
        }
    )
}

Configuration IISInstall
{
    Node $AllNodes.NodeName
    {
        WindowsFeature IIS
        {
            Name = "Web-server"
            Ensure = "Present"
        }
    }
}

IISInstall -OutputPath 'C:\Program Files\WindowsPowerShell\DscService\Configuration' -ConfigurationData
$ConfigInfoIIS

New-DSCChecksum -ConfigurationPath 'C:\Program Files\WindowsPowerShell\DscService\Configuration' -OutPath
'C:\Program Files\WindowsPowerShell\DscService\Configuration'
```

You should place the MOF files on pull server WinServer10 in a well-defined folder, typically in C:\Program Files\WindowsPowerShell\DSCService\Configuration. When the MOF file is generated, this folder location is passed as a parameter for the OutputPath attribute.

After MOF file generation, the checksum file for the configuration is generated. It maintains the integrity of configurations when they are transmitted between the pull server and LCM. The checksum file should have the same name as the MOF file with .mof.checksum as an extension. DSC provides the New-DSCChecksum cmdlet to generate the checksum file. The ConfigurationPath parameter specifies the folder location where configurations are stored. The cmdlet generates a checksum file for each configuration and saves the checksum in the folder location specified by the Output parameter.

IndexFile configuration

IndexFile is a simple configuration responsible for generating an Index.htm file on a server at the IIS default root directory (C:\Inetpub\wwwroot\). The purpose of this file is to show a "Website under maintenance" message to users. It uses a File resource and creates a file named Index.htm in the C:\Inetpub\wwwroot folder with HTML content displaying, in this example, "If you are seeing this page, it means the website is under maintenance and DSC Rocks!!!!" The configuration script that follows runs on the DemoServerWin10 server. It uses ConfigurationData to define the node name. The name is IndexFile.fcd03a8d-5a64-4982-92b3-5c89680add39, and this is the same GUID value shown earlier for the IISInstall configuration example.

```

$ConfigInfoIndex = @{
    AllNodes = @(
        @{
            NodeName = "IndexFile.fcd03a8d-5a64-4982-92b3-5c89680add39"
        }
    )
}

Configuration IndexFile
{
    Node $AllNodes.NodeName
    {
        File IndexFile
        {
            DestinationPath = "C:\inetpub\wwwroot\index.htm"
            Ensure = "Present"
            Type = "File"
            Force = $true
            Contents = "<HTML><HEAD><Title> Website under construction.</Title></HEAD><BODY> `
            <h1>If you are seeing this page, it means the website is under maintenance and DSC Rocks
            !!!!!</h1></BODY></HTML>"
        }
    }
}

IndexFile -OutputPath 'C:\Program Files\WindowsPowerShell\DscService\Configuration' -ConfigurationData
$ConfigInfoIndex
New-DSCChecksum -ConfigurationPath 'C:\Program Files\WindowsPowerShell\DscService\Configuration' -OutPath
'C:\Program Files\WindowsPowerShell\DscService\Configuration'

```

The MOF file is generated on the pull server DemoWinServer10 at the C:\Program Files\WindowsPowerShell\DSCService\Configuration folder location. The checksum file for this configuration is also generated the same way as for the previous configuration example.

Deploying the configurations

After addressing LCM and relevant authoring configurations for partial configurations, you apply them on the destination server, in this case on DemoServerWin. DSC provides the Update-DSCConfiguration cmdlet in its new release for performing consistency checks on a server. Running this cmdlet with the localhost as the ComputerName parameter simultaneously downloads from the pull server all relevant configurations (MOF content) defined by using LCM PartialConfiguration. LCM then combines all the configurations into a single MOF file and applies it to the server. The following example shows the Update-DSCConfiguration cmdlet for applying configuration:

```

PS C:\Users\me> Update-DSCConfiguration -ComputerName localhost

```

Id	Name	PSJobTypeName	State	HasMoreData	Location	Command
47	Job47	Configuratio...	Running	True	localhost	Update-DscConfiguration

Running the Get-DSCConfiguration cmdlet on this server provides all the resources (File and WindowsFeature) applied as part of DSC configuration, as shown here:

```

PS C:\Users\riteshmodi> Get-DSCConfiguration

```

```

ConfigurationName      :
DependsOn               :
ModuleName              :
ModuleVersion           :
ResourceId              :
SourceInfo              :
Credential              :
DisplayName             : Web Server (IIS)
Ensure                  : Present
IncludeAllSubFeature    : False
LogPath                 :
Name                    : Web-Server
Source                  :
PSComputerName          :

ConfigurationName      :
DependsOn               :
ModuleName              :
ModuleVersion           :

```

```

ResourceId          :
SourceInfo          :
Attributes           : {archive}
Checksum            :
Contents             :
CreatedDate          : 11/8/2014 1:40:50 PM
Credential           :
DestinationPath      : C:\inetpub\wwwroot\index.htm
Ensure              : present
Force               :
MatchSource          :
ModifiedDate         : 11/14/2014 7:09:02 AM
Recurse              :
Size                : 197
SourcePath           :
SubItems            :
Type                 : file
PSComputerName       :

```

PowershellGet and NuGet

The traditional way to install a Windows PowerShell module is to search for and find the module on the Internet and then download and install it. Windows Server 2016 Technical Preview changes the way modules are managed on a machine. It comes with the PowershellGet Windows PowerShell module built in. PowershellGet helps in finding, downloading, installing, and managing modules on a machine.

PowershellGet works with multiple providers. These providers are client tools that connect to the module repository represented by Source (location - URI). The most important provider PowershellGet works with as of Windows Server 2016 Technical Preview is NuGet. NuGet is the package manager for Windows; it provides the ability to consume not only modules but also applications and packages. NuGet can work with multiple sources, and PSGallery is the most common and preferred repository source for PowershellGet.

All the functionality for module management is included in the PowershellGet module. The first step for managing modules is to import the module into the Windows PowerShell console. Start Windows PowerShell Integrated Scripting Environment (ISE) and run the following command to load the PowershellGet module:

```

PS C:\users\me>> import-module PowershellGet -Verbose

VERBOSE: Loading module from path
'C:\Windows\system32\WindowsPowerShell\v1.0\Modules\PowershellGet\PowershellGet.psd1'.
VERBOSE: Importing function 'Find-Module'.
VERBOSE: Importing function 'Get-PSRepository'.
VERBOSE: Importing function 'Install-Module'.
VERBOSE: Importing function 'Publish-Module'.
VERBOSE: Importing function 'Register-PSRepository'.
VERBOSE: Importing function 'Set-PSRepository'.
VERBOSE: Importing function 'Unregister-PSRepository'.
VERBOSE: Importing function 'Update-Module'.
VERBOSE: Importing variable 'PSGalleryPublishUri'.
VERBOSE: Importing variable 'PSGallerySourceUri'.
VERBOSE: Importing alias 'fimo'.
VERBOSE: Importing alias 'inmo'.
VERBOSE: Importing alias 'pumo'.
VERBOSE: Importing alias 'upmo'.

```

Using the Verbose switch with the Import-Module cmdlet displays all imported functions, cmdlets, and aliases on the console. There are eight functions, two variables, and four aliases available in this module.

The first time you use the PowershellGet cmdlet, it verifies the installation of NuGet on the machine. If NuGet is not installed, a confirmation message box appears, as follows:

```

PowershellGet requires NuGet_anycpu.exe to interact with NuGet-based galleries. NuGet_anycpu.exe must be
available in 'C:\ProgramData\OneGet\ProviderAssemblies' or
'C:\Users\<username>\AppData\Local\OneGet\ProviderAssemblies'. For more information about NuGet, see
http://www.nuget.org. Do you want PowershellGet to download NuGet_anycpu.exe now?

```

Clicking Yes downloads and installs NuGet on the machine.

The cmdlets provided by PowershellGet are divided into two broad categories: modules and repository cmdlets. PowershellGet provides cmdlets for finding, installing, publishing, and updating modules from the repository. It also provides cmdlets for reading current repository settings as well as updating, registering, and unregistering them.

Two repositories are available by default: PSGallery and MSPSGallery. PowershellGet uses the NuGet provider to connect to these repositories. Running Get-PSRepository displays all existing repositories on the machine.

```
PS C:\users\me>> Get-PSRepository
```

Name	SourceLocation	OneGetProvider	InstallationPolicy
PSGallery	https://msconfiggallery.cloudapp.net/api/v2/	NuGet	Untrusted
MSPSGallery	http://www.microsoft.com/	NuGet	Trusted

Running the Get-PSRepository cmdlet with a repository name displays configurations related to that repository.

```
PS C:\users\me>> Get-PSRepository -Name PSGallery | Format-List *
```

```
Name                : PSGallery
SourceLocation       : https://msconfiggallery.cloudapp.net/api/v2/
Trusted              : False
Registered           : True
InstallationPolicy    : Untrusted
OneGetProvider        : NuGet
PublishLocation      : https://go.microsoft.com/fwlink/?LinkID=397527&clcid=0x409
ProviderOptions       : {}
```

Within the output, SourceLocation indicates the URL of the repository location, OneGetProvider identifies the package provider used to connect to the repository (NuGet in this case), Trusted indicates whether the repository is trusted, and PublishLocation shows the URL used for submission of modules (https://go.microsoft.com/fwlink/?LinkID=397527&clcid=0x409 in this case).

Running the Set-PSRepository cmdlet sets the configuration values of a repository. For example, the Set-PSRepository cmdlet that follows changes the configuration value Untrusted to Trusted for the PSGallery repository. After changing a value, you can review the new configuration by running the Get-PSRepository cmdlet.

```
PS C:\Users\me> Set-PSRepository -Name "PSGallery" -InstallationPolicy Trusted
```

```
PS C:\Users\me> Get-PSRepository
```

Name	SourceLocation	OneGetProvider	InstallationPolicy
MSPSGallery	http://www.microsoft.com/	NuGet	Trusted
PSGallery	https://www.powershellgallery.com/api/v2/	NuGet	Trusted

You use the Register-PSRepository cmdlet to add and register a new repository. The cmdlet needs the name of the repository, its source location for downloading modules, the publish location for publishing new modules to the repository, an installation policy, and the package manager to work with the repository. Running this cmdlet with Chocolatey as the name, http://chocolatey.org/api/v2/ as both source and publish location, Trusted as the installation policy value, and NuGet as the package manager Name adds a new module repository to the machine, as follows:

```
Register-PSRepository -Name "Chocolatey" -SourceLocation "http://chocolatey.org/api/v2/" `
-PublishLocation "http://chocolatey.org/api/v2/" -InstallationPolicy Trusted `
-OneGetProvider NuGet
```

After registering, the Find-Module cmdlet can search repositories, and the Install-Module cmdlet can download and install modules. Chocolatey is shown here just as an example ; it can be any repository hosting Windows PowerShell modules.

```
PS C:\Users\me> Unregister-PSRepository -Name Chocolatey
```

```
PS C:\Users\me> Find-Module
```

Running the Find-Module cmdlet with the Name parameter outputs modules related to that name. Running this cmdlet with Bing as the value for the Name parameter provides information about Bing, as shown here:

```
PS C:\Users\me> Find-Module -Name "Bing"
```

```
PS C:\Users\me> Find-Module -Name "*ing"
```

Note The Name parameter also accepts wildcard characters.

```
PS C:\Users\me> Find-Module -Name "Bing" -MinimumVersion "4.0"
```

Running the Find-Module cmdlet with Bing as the value for the Name parameter and 4.0 as the value for the RequiredVersion parameter results in an error.

150 CHAPTER 7 | Systems management

However, running the Find-Module cmdlet with Bing as the value for the Name parameter and 5.0 as the value for the RequiredVersion parameter results in output with details about the Bing module.

```
PS C:\Users\me> Find-Module -Name "Bing" -RequiredVersion "5.0"
```

Repository	Version	Name	Description
PSGallery	5.0	Bing	A few functions for working with the new Bing APIs

After finding the relevant modules, the next step is to install the module. PowershellGet provides the Install-Module cmdlet specifically for this purpose. This cmdlet is very similar to Find-Module. It also takes Name, RequiredVersion, and MinimumVersion as parameters.

The code block that follows demonstrates running Install-Module with the Name parameter and the Verbose switch. You can use MinimumVersion or RequiredVersion along with the Name parameter. Notice that the last line in the output suggests that module is installed. Also, note that the modules are downloaded by default at the \$env:ProgramFiles\WindowsPowerShell\Modules folder location. Windows PowerShell uses this folder location to install modules.

```
PS C:\Users\me> Install-Module -Name bing -Verbose
VERBOSE: In PSModule Provider - 'Get-DynamicOptions'.
VERBOSE: In PSModule Provider - 'Get-DynamicOptions'.
VERBOSE: In PSModule Provider - 'Get-DynamicOptions'.
VERBOSE: In PSModule Provider - 'Get-DynamicOptions'.
VERBOSE: In PSModule Provider - 'Find-Package'.
VERBOSE: OPTION: MessageResolver => Microsoft.PowerShell.OneGet.CmdLets.GetMessageString
VERBOSE: OPTION: ProviderName => PSModule
VERBOSE: OPTION: Verbose => True
VERBOSE: OPTION: Name => bing
VERBOSE: The -Repository parameter was not specified. PowerShellGet will use all of the registered
repositories.
VERBOSE: Getting the provider object for the OneGet Provider 'NuGet'.
VERBOSE: The specified Location is 'http://www.microsoft.com/' and OneGetProvider is 'NuGet'.
VERBOSE: Calling 'NuGet::FindPackage'
VERBOSE: Calling 'CommonServiceProvider::GetKnownFolder'
VERBOSE: Getting the provider object for the OneGet Provider 'NuGet'.
VERBOSE: The specified Location is 'https://www.powershellgallery.com/api/v2/' and OneGetProvider is 'NuGet'.
VERBOSE: Calling 'NuGet::FindPackage'
VERBOSE: Calling 'CommonServiceProvider::GetKnownFolder'
VERBOSE: Performing the operation "" on target "Version '5.0' of module 'Bing'".
VERBOSE: In PSModule Provider - 'Install-Package'.
VERBOSE: The FastPackageReference is 'NuGet|Bing|5.0|https://www.powershellgallery.com/api/v2/'.
VERBOSE: OPTION: MessageResolver => Microsoft.PowerShell.OneGet.CmdLets.GetMessageString
VERBOSE: OPTION: ProviderName => PSModule
VERBOSE: OPTION: Verbose => True
VERBOSE: OPTION: Name => bing
VERBOSE: Version '5.0' of module 'Bing' is already installed at 'C:\Program
Files\WindowsPowerShell\Modules\Bing'.
```

At this point, you can use the Bing module by importing it into the current Windows PowerShell runspace by using the Import-Module cmdlet. After the initial installation, running Update-Module updates the existing modules. This module takes the Name and RequiredVersion parameters, but it does not take the MinimumVersion parameter, as shown here:

```
PS C:\Users\me> Update-Module -Name Bing
PS C:\Users\me> Update-Module -Name "Bing" -RequiredVersion "5.0"
```

There is also a Publish-Module cmdlet for adding newer modules to the repository.

System Center 2016

By John McCabe

Just like Windows Server, System Center gets an updated edition, too. In this section, we will detail what is new for System Center 2016. The core focus of System Center 2016 is on hybrid management—how can we manage the cloud natively from System Center, but also how can we use the cloud to extend the functionality of system center or manage the environment from the cloud.

Operations Management Suite and Intune are the management functions within the cloud that complements the System Center 2016 suite.

As you can imagine, a lot of what is new in System Center 2016 also focuses on ensuring that we can support the new capabilities in Windows Server 2016 Technical Review. Also, System Center 2016 is designed to truly facilitate the software-defined datacenter (SDDC) and gives you all the tools you require to accomplish this. In the following table, we give you a breakdown of some of the new features coming in System Center 2016 with respect to their general management areas:

Focus area	Features
Device management	Windows 10 deployment support MDM enrollment with Azure Active Directory Access restriction based on device enrollment and policy
Provisioning	Support for Windows Server 2016 Technical Review Hyper-V features Rolling cluster upgrades Simplified networking Shielded virtual machine (VM) provisioning Guarded host management VMWare vCenter 5.5 support
Monitoring	Nano Server Windows storage SMS-S support MP catalog improvements Performance improvements Enhanced data visualization Improved Linux support Improved network support
Automation	Migration to cloud SCO integration packs and runbooks SMA support native Windows PowerShell Windows Management Framework 5.0 PowerShell ISE plug-in support for SMA runbooks
Self-service	Improved usability and performance HTML 5 self-service portal New Microsoft Exchange connector
Data protection	Azure Express Route supported Shielded VM support Storage spaces direct

Traditionally, System Center was geared toward managing your on-premises infrastructure. This continued to evolve in the previous version, and is yet even more of a focus in System Center 2016.

You can use System Center 2016 to manage your Cloud Environments, as well. For example, do you want to know the health of your Office 365 Subscription? In System Center 2016, you now can gather this information. Figure 7-1 shows a sample of the dashboard available in the Management Pack for Microsoft Office 365.




Figure 7-1: Office 365 dashboard view

The Subscription Health shown in Figure 7-1 validates whether your configuration is correct and you have a successful connection to Office 365. You would configure a set of credentials as part of deploying this management pack, which would need the appropriate permissions to connect to the subscription. Figure 7-1 also illustrates the areas of the Office 365 subscription that have active incidents. The active incidents, much like other Operations Managers alerts, have health state information and a knowledge center with possible resolution steps attached.

The previous example shows only Office 365, but it is only a small part of the capabilities inherent in System Center 2016 for managing a public cloud service. You also can extend management into public cloud environments from the perspectives of Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

Let's look at a quick scenario to show you how you can use the entire suite across all environments. Contoso Limited is a company that produces a simple widget. The widget is such a useful product that demand has grown exponentially since its release, and now we need a fully Agile IT solution that can meet the needs and demands of not only the Contoso employees but also the Contoso customers.

Contoso customers want to be able to buy this widget at any time, from any place. And, yes, of course from any device! What does this mean in real terms? Simply put, Contoso needs a system that will be available 24x7. What does it mean in terms of infrastructure and management? First, we have two websites connecting to database servers hosted in the cloud. The first website is for customers, and the second is for the sales staff of Contoso to check inventory, look at orders, and so on. We also have mobile services so that phone apps can connect any time and place orders. The website for the sales staff uses Azure Active Directory to authenticate users. This in turn means Contoso has extended its Active Directory to Azure to accommodate this. Contoso also uses Office 365 for its email and collaboration. Finally, the company integrated its telephony system with Office365 to allow for Global telecommunications.

Contoso is stretched across public and private cloud. It has a global customer base that is supported by a global employee base to serve those customer needs. Using System Center 2016, the company can provision infrastructure when required, manage the application estate, integrate development and operations together so that true metrics showing how an application is being used can be reflected all the way in the development chain, provide detailed remediation and diagnostic tasks, and so on.

Although this might not seem like an example of real infrastructure today, the modern enterprise is evolving. Customers can be anywhere, and if you come from an environment where you host your systems based out of a region that is far away from a customer base, their end-user experience will not be good. Suppose that you host your central datacenters in the United States and you have offices and customers in India. The user experience will be bad because of the inherent distance-based latency alone. Let's further suppose that you decide to not use the cloud to host your applications and instead decide to open a local datacenter in India. What happens then is you duplicate your management structure and the system becomes more complex to manage. There are multiple scenarios we could discuss which highlight the evolving nature of IT and how to manage it.

When designing System Center 2016, a key focus was to address the Contoso scenario but not just from managing on-premises out to the cloud, but in either direction. Another key area was how to keep up with cloud cadence, the investment to keep System Center 2016 up to date and move at the speed of the cloud is considerable, but making investments to have a hybrid relationship between a management solution based in the cloud and on-premises investments make sense. In the next section, we will look at Operations Management Suite which can help us take this journey

Operations Management Suite

By John McCabe

Operations Management Suite (OMS) is a cloud-based management solution that can complement on-premises investments in System Center 2016 or stand independently. Figure 7-2 shows how you can manage multivendor environments and get the best experience possible by combining System Center 2016 and OMS to manage your hybrid IT world.




Figure 7-2: Using System Center 2016 and OMS to manage a multivendor environment

Before diving into how we can take advantage of the capabilities of OMS and System Center 2016, let's describe what's available in OMS. OMS can be divided into the following four primary areas:

Focus area	Description
Log analytics	Search for patterns, identify problems across a multitude of different log sources and provide real-time insights into what is happening in your environment. Integrate into Microsoft PowerBI dashboards for powerful visualizations.
IT automation	Automate simple and complex tasks in your IT environment; directly integrate with applications and provide source control for your automation environment; connect and manage resources across datacenters.
Backup and recovery	Back up your workloads directly to the cloud and use the cloud as a recovery point. Alternatively, replicate your workloads from VMware or Hyper-V and use the cloud as a recovery site.

Security and compliance	Continually assess and understand what is happening in your environment, from who is signing in, to a new risk that is highlighted in your environment.
-------------------------	---

The key takeaway here is the ability to be hybrid. This is particularly relevant if you have made a large investment on-premises and want to use OMS and its features along with System Center 2016. Even if you haven't made any investment into System Center on-premises but like what OMS can offer, no problem: You can take advantage of OMS to manage your existing cloud or on-premises estates.

To begin, whether you have OMS deployed or not, you must create an OMS workspace. To do so, sign in to <https://portal.azure.com>, and then click New. Next, type **Operational Insights**, click Operational Insights (see Figure 7-3), and then click Create.




Figure 7-3: Creating your OMS workspace

This starts the service management portal. Populate the settings to match those shown in Figure 7-4 and then click Create Workspace. You have a choice of tier; for most users, the Free tier is a great way to explore the power and benefits of OMS.




Figure 7-4: Creating your OMS workspace, part 2

Note Currently, Operational Insights is being updated to the Azure Resource Manager experience end to end. It is expected to be released in Calendar 2016.

After the workspace is created, its status is listed as Active. At this point, highlight the workspace, and then, on the bottom toolbar, click Manage. (You also can click the small arrow on the name of the workspace.) This brings you to some basic settings, one of which you might want to consider

implementing. In Azure, a lot of services have the ability to write log files directly to a Storage account. You can add this account to the workspace so that you can later perform analysis on it.

When you sign in to the workspace, the first thing you need to do is click the Get Started tile, as shown in Figure 7-5.




Figure 7-5: The main workspace

There are three main tasks to accomplish when getting started with OMS. When you click Get Started, a wizard-like experience guides you through the process of selecting solutions. Solutions are like management packs in OMS. They contain all of the intelligence and rules against which machines in the environment you present will be assessed. Solutions are updated on a cloud cadence, and new solutions are continually being developed and added to the overall portfolio based on customer demands and requirements.

Figure 7-6 depicts the first step for configuring OMS. To get up and running requires that you select some solutions. In the pane on the left, click Add Solutions. These solutions won't really do anything until they have machines to work against, so you can technically select all of them or only the ones with which you are interested in working.




Figure 7-6: Step one: selecting solutions

Next, click Connect A Data Source. Figure 7-7 shows a mixture of options from which you can select.




Figure 7-7: Step two: connecting data sources

Here, you have three basic questions to answer:

1. Do you want to deploy an agent directly to a machine and register directly with OMS?
2. Do you want to connect an operations manager deployment to OMS?
3. Do you want to add a Storage account that contains log data?

Your answers will determine which steps you take to complete the installation. If you want the destination machine to report directly to OMS, download the agent and install it on the machine. During the installation, you will be prompted to select the type of deployment that you want to

register the agent against. The agent itself is the Microsoft Management Agent, which can be registered directly with OMS or an OMS server, as shown in Figure 7-8.




Figure 7-8: Installing Microsoft Monitoring Agent

When you select the Connect The Agent To Microsoft Azure Operational Insights check box, you are prompted for the workspace ID and key. You can obtain these from the Operational Workspace, as previously shown in Figure 7-7, and enter them into the boxes, as shown in Figure 7-9.




Figure 7-9: Configuring Workspace ID and Key

The agent will complete its installation and then register with the OMS workspace. When the agent has registered with OMS, you will see a green check mark beside Step 2, and you will see one server connected in the OMS workspace.

Finally, you can configure Step 3, Add Logs. Figure 7-10 shows the different log types you can select. For example, in the search box, you can type **free** for Windows Event Logs and type **System** and you will see it will try to resolve the available Logs. Ensure that you click Save.




Figure 7-10: Adding logs

From here the rules are downloaded to the agent, as normal, and processed. Data will be uploaded to the portal and assessed. The main solution gallery will be updated with the latest information pulled from the system. You can add additional solutions from the solutions gallery when you need them.

Figure 7-11 presents an updated dashboard after information has been uploaded.




Figure 7-11: Updated dashboard

You can click each site to view more information. From here you can also configure additional items such as Automation, Backup, and Azure Site Recovery. You can use all three areas in hybrid scenarios to manage cloud resources and on-premises resources from the cloud.

From here, you can explore Log Search and all additional solutions, as shown in Figure 7-12.

 AD Assessment Free Assess the risk and health of Active Directory environments.	 Alert Management Free Manage your Operations Manager alerts across your servers.	 Azure Networking Analytics Coming Soon Gain insight into your Azure Network data.	 Containers Coming Soon See Docker container performance metrics and logs from containers across your public or private cloud environments.	 Network Performance Monitor Coming Soon Offers near real time monitoring of network performance parameters like loss and latency.	 Security and Audit Free Provides the ability to explore security related data and helps identify security breaches.	 SQL Assessment Free Assess the risk and health of SQL Server environments.	 Wire Data Coming Soon Provides the ability to explore wire data and helps identify network related issues.	 Automation Owned Automate time consuming and frequently repeated tasks in the cloud and on premises.	 Change Tracking Owned Track configuration changes across your servers.	 System Update Assessment Owned Identify missing system updates across your servers.
 AD Replication Status Free Identify Active Directory replication issues in your environment.	 App Dependency Monitor Coming Soon Automatically discover and map servers and their dependencies in real time.	 Configuration Assessment Free Identify configuration problems across your servers.	 Key Vault Coming Soon Understand your Key Vault usage through Analysis of Key Vault logs.	 Office 365 Coming Soon Get full visibility into your Office 365 user activities, perform forensics as well as audit and compliance.	 Service Fabric Coming Soon Identify and troubleshoot issues across your Service Fabric cluster.	 Surface Hub Free Provides the ability to monitor Microsoft Surface Hub devices.	 Malware Assessment Owned View status of antivirus and antimalware scans across your servers.	 Backup Owned Manage Azure IaaS VM backup and Windows Server backup status for your backup vault.	 Azure Site Recovery Owned Monitor virtual machine replication status for your Azure Site Recovery vault.	

Figure 7-12: Solution gallery in OMS

About the author



John McCabe works for Microsoft as a senior premier field engineer. In this role, he has worked with the largest customers around the world, supporting and implementing cutting-edge solutions on Microsoft Technologies. In this role, he is responsible for developing core services for the Enterprise Services Teams. John has been a contributing author to several books, including *Mastering Windows Server 2012 R2* from Sybex, *Mastering Lync 2013* from Sybex, and *Introducing Microsoft System Center 2012* from Microsoft Press.

John has spoken at many conferences around Europe, including TechEd and TechReady. Prior to joining Microsoft, John was an MVP in Unified Communications with 15 years of consulting experience across many different technologies such as networking, security, and architecture.

About the contributors

Shabbir Ahmed is a partner enterprise architect with the PEAT team at Microsoft. He specializes in designing solutions for partners and customers by linking and applying complex technologies to their business strategies. Shabbir is a creative thinker with high energy and enthusiasm. He studied Executive PG in IT at the Indian Institute of Management (IIM-K) and holds many awards and certifications, including CCIE#21327, MVP, ISO 27001 LA, CEH, MCT, IPv6 forum (Gold). You can find his LinkedIn profile at <http://in.linkedin.com/pub/shabbir-ahmed/58/575/209>.

David Branscome has worked at Microsoft for six years, in both consulting and premier support. Currently, he is a premier field engineer, supporting federal customers deploying and managing Microsoft technologies.

Yuri Diogenes is a senior content developer on the Microsoft CSI Enterprise Mobility Team. He has a Master of Science degree in Cybersecurity from UTICA College, is a Certified Information Systems Security Professional (CISSP). Yuri has authored many IT pro and security books from Microsoft Press and Syngress. You can follow him on Twitter @yuridiogenes and on his blog at <http://blogs.technet.com/yuridiogenes/>.

Matt Garson is a program manager in the Microsoft Operating Systems Group, working on the storage and file systems team. Matt has been a member of the storage team since joining Microsoft, working on both file systems and software-defined storage technologies.

Claus Joergensen is a principal program manager on the Microsoft Windows Server HA and Storage team. His current responsibilities include high availability, scale out, and shared-nothing storage. His previous responsibilities included Scale-Out File Server, specifically SMB Transparent Failover, SMB Scale-Out, and VSS for SMB File Shares. Prior to this role, Claus held a number of positions at Microsoft, including principal consultant in Microsoft Services.

John Marlin has been with Microsoft since April 1992 and was in the original Cluster Group that started with NT 4. John delivers high-availability (Cluster/Hyper-V) training to internal groups as well as to external partners, and he has delivered Failover Clustering sessions at multiple TechEd conferences. He has written articles for both Windows IT Magazine and TechNet Magazine and is a regular contributor to the AskCore Blog Site. He has also contributed to the Microsoft Press books *Introducing Windows Server 2012 R2* and *Optimizing and Troubleshooting Hyper-V Storage*. As a part of his work with the current support group, John deals with clustering, storage, and virtualization. He also handles interoperability issues with networking and Active Directory. You can follow John's blog at <http://blogs.technet.com/b/askcore/archive/tags/john+marlin/>.

Andrew Mason is a principal PM manager in the Enterprise and Cloud Division at Microsoft, where he works on Windows Server and Services. Andrew's team is responsible for the Windows Server platform and tools, including the various installation options of Windows Server, such as Server Core. He has been involved with Windows Server in one way or another at Microsoft since before Windows NT 3.1.

Meir Mendelovich is a principle program manager working on security products in Microsoft Enterprise & Cloud Division. He worked on OMS Security, Windows Server Web Application Proxy, Azure AD Application Proxy, and the legacy products in this area (UAG and TMG).

Robert Mitchell is a 20-plus year employee of Microsoft. He occasionally writes for Windows IT Pro magazine (see <http://windowsitpro.com/author/robert-mitchell>) and runs a public-facing TechNet blog at http://blogs.technet.com/b/tip_of_the_day/. Robert enjoys reading, video games, and spending time with his lovely wife, Heather.

Ritesh Modi is an architect with Microsoft Services. He has been with Microsoft for the past four years, working on datacenter and Azure infrastructure, DevOps automation, and Managed Services. He's an expert on Azure, DevOps, Windows PowerShell, SharePoint, SQL Server, and System Center. He has been responsible for defining the Managed Services offering and its platform. He is also involved in creating automation library and cloudburst implementation. He has spoken at multiple conferences, including TechEd and PowerShell Asia conference, he performs internal training, and he is a published author for MSDN magazine. Ritesh has more than a decade of experience in building and deploying enterprise solutions for customers. You can read his blog at <http://automationnext.wordpress.com> and follow him on Twitter @automationnext.

Ned Pyle is a senior program manager in the Microsoft Windows Server engineering group, managing replication and remote file protocols. His previous role was a technical lead within Microsoft escalation support, where he was the founder and editor of the infamous AskDS blog. Prior to joining Microsoft, he spent eight years in IT consulting and was a United States Marine infantryman. He lives in Seattle, Washington, with his wife and their dogs.

Colin Robinson is a Program Manager in the Cloud and Enterprise division with 17 years of experience, with the last 4 years of that at Microsoft. He and his wife previously lived in the Netherlands and now enjoy traveling the United States and the world together. He has been to more than 1,000 concerts and enjoys martial arts and photography as hobbies.

John Savill is a principal solutions architect for which he focuses on infrastructure solutions. John has worked with Windows since NT 3.1 and has authored multiple books on Windows, Hyper-V, and Azure. You can follow John on Twitter @NTFAQGuy and on his blog <http://www.savilltech.com/blog>.

Ramnish Singh has been recognized as a leader in the field of IT architecture by both Microsoft and the International Association of Software Architects (IASA). He earned the title of Microsoft Certified Architect (Infrastructure) from Microsoft and Certified IT Architect (Solutions) from IASA. He serves on the boards of Microsoft Certified Architect and International Association of Software Architect Programs. He has authored many publications for *Architect Journal*, MSDN, TechNet, and CSI.

Deepak Srivastava is a partner consultant with Microsoft and has experience in architecture, design, and migration projects at financial, legal, IT, and Fortune 500 organizations, with deep technical expertise in datacenter and cloud infrastructure technologies, including Microsoft Azure, Windows Server, Hyper-V, System Center, and VMware vCloud Suite. He has presented at various cloud events worldwide and has contributed to several books. In his personal time, he prefers to play table tennis, read books, and watch movies. For more information about Deepak, go to <https://www.linkedin.com/in/deepakksrivastava>.



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